

AD-A162 344


HEURISTIC CLASSIFICATION(U) STANFORD UNIV CA DEPT OF
COMPUTER SCIENCE W J CLANCEY AUG 85 STAN-CS-85-1066
N00014-85-K-0305

1/1

UNCLASSIFIED

F/G 9/2


NL



END

END

END



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

June 1985
AUGUST

12

Report No. STAN-CS-85-1066
Also numbered KSL-85-5

AD-A162 344

Heuristic Classification

by


William J. Clancey

Department of Computer Science

Stanford University
Stanford, CA 94305

DTIC
FLECTE
DEC 17 1985
S **D**
B

FILE COPY



DISTRIBUTION STATEMENT A
Approved for public release
Distribution Unlimited

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER Tech Report #12	2. GOVT ACCESSION NO. AD-A162344	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Heuristic Classification	5. TYPE OF REPORT & PERIOD COVERED Technical, June 1985	
7. AUTHOR(s) William J. Clancey	6. PERFORMING ORG. REPORT NUMBER	
9. PERFORMING ORGANIZATION NAME AND ADDRESS Stanford Knowledge Systems Laboratory 701 Welch Road, Bldg. C Department of Computer Science Palo Alto, CA 94304	8. CONTRACT OR GRANT NUMBER(s) N00014-85K-0305	
11. CONTROLLING OFFICE NAME AND ADDRESS Personnel and Training Research Programs Office of Naval Research (Code 458) Arlington, VA 22217	10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NR702-003	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) ONR Representative - Mr. Robin Simpson Durand Aeronautics Building, Rm. 165 Stanford University, Stanford, CA 94305	12. REPORT DATE August 1985	
	13. NUMBER OF PAGES	
	15. SECURITY CLASS. (of this report) Unclassified	
	15a. DECLASSIFICATION/DOWNGRADING SCHEDULE	
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release: distribution unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES Also KSL Memo 85-5 to appear in <u>Journal of Artificial Intelligence</u> .		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) A broad range of well-structured problems--embracing forms of diagnosis, catalog selection, and skeletal planning--are solved in "expert systems" by the method of heuristic classification. These programs have a characteristic inference structure that systematically relates data to a pre-enumerated set of solutions by abstraction, heuristic association, and refinement. In contrast with previous descriptions of classification reasoning, particularly in psychology, this analysis emphasizes the role of a heuristic in routine problem solving as a non-hierarchical, direct association between concepts. In contrast with other		

19. KEY WORDS (Continued)

20 ABSTRACT (Continued)

descriptions of expert systems, this analysis specifies the knowledge needed to solve a problem, independent of its representation in a particular computer language. The heuristic classification problem-solving model provides a useful framework for characterizing kinds of problems, for designing representation tools, and for understanding non-classification (constructive) problem-solving methods.

HEURISTIC CLASSIFICATION

William J. Clancey

Stanford Knowledge Systems Laboratory
Department of Computer Science
701 Welch Road, Building C
Palo Alto, CA 94304

The studies reported here were supported (in part) by:

The Office of Naval Research
Personnel and Training Research Programs.
Psychological Sciences Division.
Contract No. N00014-85K-0305

The Josiah Macy, Jr. Foundation
Grant No. B852005
New York City

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Office of Naval Research or the U.S. Government.

Approved for public release; distribution unlimited. Reproduction in whole or in part is permitted for any purpose of the United States Government.

Table of Contents

Abstract	1
1. INTRODUCTION	1
2. THE HEURISTIC CLASSIFICATION METHOD DEFINED	4
2.1. Simple classification	4
2.2. Data abstraction	5
2.3. Heuristic classification	6
3. EXAMPLES OF HEURISTIC CLASSIFICATION	8
3.1. SACON	9
3.2. GRUNDY	9
3.3. SOPHIE III	12
4. UNDERSTANDING HEURISTIC CLASSIFICATION	12
4.1. Schemas vs. definitions	14
4.2. Alternative encodings of schemas	15
4.3. Relating heuristics to conceptual graphs	18
4.4. Relating inference structure to conceptual graphs	21
4.5. Pragmatics of defining concepts	24
5. ANALYSIS OF PROBLEM TYPES IN TERMS OF SYSTEMS	27
5.1. What gets selected?	27
5.2. Background: Problem categories	28
5.3. A system-oriented approach	29
5.4. Configuration and planning	32
5.5. Combinations of system problems	36
6. INFERENCE STRATEGIES FOR HEURISTIC CLASSIFICATION	39
6.1. Focusing in heuristic classification	39
6.2. Causal-process classification	43
6.2.1. Electronic and medical diagnosis compared	44
6.2.2. Inference control for coherency	46
6.2.3. Multiple solutions and levels of detail	48
6.2.4. Constructing implication paths vs. constructing solutions	48
7. CONSTRUCTIVE PROBLEM SOLVING, AN INTRODUCTION	49
8. RELATING TOOLS, METHODS, AND TASKS	52
9. KNOWLEDGE-LEVEL ANALYSIS	55
10. RELATED ANALYSES IN PSYCHOLOGY AND ARTIFICIAL INTELLIGENCE	56
11. SUMMARY OF KEY OBSERVATIONS	59
12. IMPLICATIONS	60
Acknowledgments	61



✓
PER CALL JC

A-1

List of Figures

Figure 2-1:	Mycin's classification of bacteria	5
Figure 2-2:	Inference structure of MYCIN	7
Figure 2-3:	Inference structure of heuristic classification	8
Figure 3-1:	Inference structure of SACON	10
Figure 3-2:	Inference structure of GRUNDY	11
Figure 3-3:	Inference structure of SOPHIE	13
Figure 4-1:	Schema describing the concept CLUSTER-HEADACHE and some related concepts	16
Figure 4-2:	A heuristic rule expanded as a conceptual graph	19
Figure 4-3:	Using a general rule to work backwards from a solution	20
Figure 4-4:	Conceptual relations used in heuristic classification	21
Figure 4-5:	Typical conceptual relations in simplest form of heuristic classification	23
Figure 4-6:	Conceptual graph of the term "brain-mass-lesion"	25
Figure 5-1:	Generic operations for synthesizing a system	30
Figure 5-2:	Generic operations for analyzing a system	30
Figure 5-3:	The design problem seen from two perspectives	33
Figure 5-4:	Simple FLIPP display, encoding rules $A \rightarrow B$, $A \rightarrow C$, and $D \rightarrow C$. (From Cox, 1984.)	36
Figure 6-1:	Backward deduction in NEOMYCIN to confirm a solution	42
Figure 6-2:	Inference structure of causal process classification	44
Figure 6-3:	Interacting data in classification	46
Figure 6-4:	Multiple explanations for a datum	47

To understand something as a specific instance of a more general case-- which is what understanding a more fundamental principle or structure means--is to have learned not only a specific thing but also a model for understanding other things like it that one may encounter. (Bruner, 1960)

Abstract

A broad range of well-structured problems—embracing forms of diagnosis, catalog selection, and skeletal planning—are solved in *expert systems* by the method of *heuristic classification*. These programs have a characteristic inference structure that systematically relates data to a pre-enumerated set of solutions by abstraction, heuristic association, and refinement. In contrast with previous descriptions of classification reasoning, particularly in psychology, this analysis emphasizes the role of a heuristic in routine problem solving as a *non-hierarchical, direct association between concepts*. In contrast with other descriptions of expert systems, this analysis specifies the knowledge needed to solve a problem, independent of its representation in a particular computer language. The heuristic classification problem-solving model provides a useful framework for characterizing kinds of problems, for designing representation tools, and for understanding non-classification (constructive) problem-solving methods.

1. INTRODUCTION

Over the past decade, a variety of heuristic programs, commonly called "expert systems," have been written to solve problems in diverse areas of science, engineering, business, and medicine. Developing these programs involves satisfying an interacting set of requirements: Selecting the application area and specific problem to be solved, bounding the problem so that it is computationally and financially tractable, and implementing a prototype program—to name a few obvious concerns. With continued experience, a number of programming environments or "tools" have been developed and successfully used to implement prototype programs (Hayes-Roth, et al., 1983). Importantly, the representational units of tools (such as "rules" and "attributes") provide an orientation for identifying manageable subproblems and organizing problem analysis. Selecting appropriate applications now often takes the form of relating candidate problems to known computational methods, our tools.

Yet, in spite of this experience, when presented with a given "knowledge engineering tool," such as EMYCIN (van Melle, 1979), we are still hard-pressed to say what kinds of problems it can be used to solve well. Various studies have demonstrated advantages of using one representation language instead of another—for ease in specifying knowledge relationships, control of reasoning, and perspicuity for maintenance and explanation (Swartout, 1981, Aiello, 1983, Aikins, 1983, Clancey, 1983a, Clancey and Letsinger, 1984). Other studies have characterized in low-level terms why a given problem might be inappropriate for a given language, for example, because data are time-varying or subproblems interact (Hayes-Roth, et al., 1983). While these studies reveal the weaknesses and limitations of the rule-based

formalism, in particular, they do not clarify the form of analysis and problem decomposition that has been so successfully used in these programs. In short, attempts to describe a mapping between *kinds of problems* and programming languages have not been satisfactory because they don't describe what a given program *knows*: Applications-oriented descriptions like "diagnosis" are too general (e.g., solving a diagnostic problem doesn't necessarily require a device model), and technological terms like "rule-based" don't describe what kind of problem is being solved (Hayes, 1977, Hayes, 1979). We need a better description of what heuristic programs do and know—a computational characterization of their competence—independent of task and independent of programming language implementation. Logic has been suggested as a basis for a "knowledge-level" analysis to specify what a heuristic program does and might know (Nilsson, 1981, Newell, 1982). However, we have lacked a set of terms and relations for doing this.

In an attempt to characterize the *knowledge-level competence* of a variety of expert systems, a number of programs were analyzed in detail.¹ There is a striking pattern: These programs proceed through easily identifiable phases of data abstraction, heuristic mapping onto a hierarchy of pre-enumerated solutions, and refinement within this hierarchy. In short, these programs do what is commonly called *classification*, but with the important twist of *relating concepts in different classification hierarchies by non-hierarchical, uncertain inferences*. We call this combination of reasoning *heuristic classification*.

Note carefully: The heuristic classification model characterizes *a form of knowledge and reasoning*—patterns of familiar problem situations and solutions, heuristically related. In capturing problem situations that tend to occur and solutions that tend to work, this knowledge is essentially *experiential*, with an overall form that is problem-area independent. Heuristic classification is a method of computation, not a kind of problem to be solved. Thus, we refer to "the heuristic classification method," not "classification problem."

Focusing on epistemological content rather than representational notation, this paper proposes a set of terms and relations for describing the knowledge used to solve a problem by the heuristic classification method. Subsequent sections describe and illustrate the model in the analysis of MYCIN, SACON, GRUNDY, and SOPHIE III. Significantly, a knowledge-level description of these programs corresponds very well to psychological models of expert problem solving. This suggests that the heuristic classification problem-solving model captures general principles of how experiential knowledge is organized and used, and thus generalizes some cognitive science results. A thorough discussion relates the model to schema research; and use

¹Including: Ten rule-based systems [MYCIN, PUFF, CLOT, HEADMED, SACON from the EMYCIN family (Buchanan and Shortliffe, 1984), plus WINE, BANKER, The Drilling Advisor, and other proprietary systems developed at Teknowledge, Inc.], a frame-based system (GRUNDY), and a program coded directly in LISP (SOPHIE III).

of a conceptual graph notation shows how the inference-structure diagram characteristic of heuristic classification can be derived from some simple assumptions about how data and solutions are typically related (Section 4). Another detailed discussion then considers "what gets selected," possible kinds of solutions (e.g., diagnoses). A taxonomy of problem types is proposed that characterizes *solutions of problems* in terms of synthesis or analysis of some *system* in the world (Section 5). We finally turn to the issue of inference control in order to further characterize tool requirements for heuristic classification (Section 6), segueing into a brief description of constructive problem solving (Section 7).

This paper explores different perspectives for describing expert systems; it is not a conventional description of a particular program or programming language. The analysis does produce some specific and obviously useful results, such as a distinction between electronic and medical diagnosis programs (Section 6.2). But there are also a few essays with less immediate payoffs, such as the analysis of problem types in terms of systems (Section 5) and the discussion of the pragmatics of defining concepts (Section 4.5). Also, readers who specialize in problems of knowledge representation should keep in mind that the discussion of schemas (Section 4) is an attempt to clarify the knowledge represented in rule-based expert systems, rather than to introduce new representational ideas.

From another perspective, this paper presents a methodology for analyzing problems, preparatory to building an expert system. It introduces an intermediate level of knowledge specification, more abstract than specific concepts and relations, but still independent of implementation language. Indeed, one aim is to afford a level of awareness for describing expert system design that enables knowledge representation languages to be chosen and used more deliberately.

We begin with the motivation of wanting to formalize what we have learned about building expert systems. How can we classify problems? How can we select problems that are appropriate for our tools? How can we improve our tools? Our study reveals patterns in knowledge bases: Inference chains are not arbitrary sequences of implications, they compose relations among concepts in a systematic way. Intuitively, we believe that understanding these high-level knowledge structures, implicitly encoded in today's expert systems, will enable us to teach people how to use representation languages more effectively, and also enable us to design better languages. Moreover, it is a well-established principle for designing these programs that the knowledge people are trying to express should be stated explicitly, so it will be accessible to auxiliary programs for explanation, teaching, and knowledge acquisition (e.g., (Davis, 1976)).

Briefly, our methodology for specifying the knowledge contained in an expert system is based on:

- a computational distinction between selection and construction of solutions;

- a relational breakdown of concepts, distinguishing between abstraction and heuristic association and between subtype and cause, thus revealing the classification nature of inference chains; and
- a categorization of problems in terms of synthesis and analysis of *systems in the world*, allowing us to characterize inference in terms of a sequence of classifications involving some system.

The main result of the study is the model of heuristic classification, which turns out to be a common problem-solving method in expert systems. Identifying this computational method is not to be confused with advocating its use. Instead, by giving it a name and characterizing it, we open the way to describing when it is applicable, contrasting it with alternative methods, and deliberately using it again when appropriate.

As one demonstration of the value of the model, classification in well-known medical and electronic diagnosis programs is described in some detail, contrasting different perspectives on what constitutes a diagnostic solution and different methods for *controlling inference* to derive *coherent* solutions. Indeed, an early motivation for this study was to understand how NEOMYCIN, a medical diagnostic program, could be generalized. The resulting tool, called HERACLES (roughly standing for "Heuristic Classification Shell") is described briefly, with a critique of its capabilities in terms of the larger model that has emerged.

In the final sections of the paper, we reflect on the adequacy of current knowledge engineering tools, the nature of a knowledge-level analysis, and related research in psychology and artificial intelligence. There are several strong implications for the practice of building expert systems, designing new tools, and continued research in this field. Yet to be delivered, but promised by the model, are explanation and teaching programs tailored to the heuristic classification model, better knowledge acquisition programs, and demonstration that thinking in terms of heuristic classification makes it easier to choose problems and build new expert systems.

2. THE HEURISTIC CLASSIFICATION METHOD DEFINED

We develop the idea of the heuristic classification method by starting with the common sense notion of classification and relating it to the reasoning that occurs in heuristic programs.

2.1. Simple classification

As the name suggests, the simplest kind of classification is identifying some unknown object or phenomenon as a member of a known class of objects, events, or processes. Typically, these classes are stereotypes that are hierarchically organized, and the process of identification is one of matching observations of an unknown entity against features of known classes. A

paradigmatic example is identification of a plant or animal, using a guidebook of features, such as coloration, structure, and size. MYCIN solves the problem of identifying an unknown organism from laboratory cultures by matching culture information against a hierarchy of bacteria (Figure 2-1).²

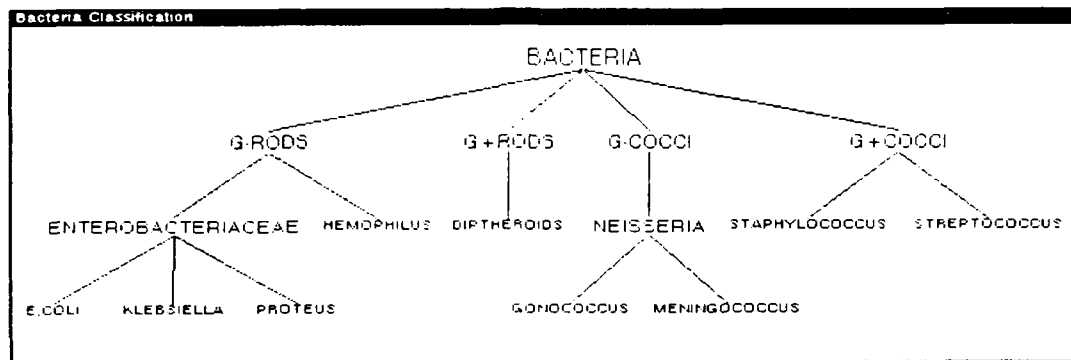


Figure 2-1: Mycin's classification of bacteria

The essential characteristic of classification is that the problem solver *selects* from a set of pre-enumerated solutions. This does not mean, of course, that the "right answer" is necessarily one of these solutions, just that the problem solver will only attempt to match the data against the known solutions, rather than construct a new one. Evidence can be uncertain and matches partial, so the output might be a ranked list of hypotheses. Besides matching, there are several *rules of inference* for making assertions about solutions. For example, evidence for a class is indirect evidence that one of its subtypes is present.

2.2. Data abstraction

In the simplest problems, data are solution features, so the matching process is direct. For example, an unknown organism in MYCIN can be classified directly given the supplied data of Gram stain and morphology. The features "Gram-stain negative" and "rod-shaped" match a class of organisms. The solution might be refined by getting information that allows subtypes to be discriminated.

For many problems, solution features are not supplied as data, but are inferred by *data abstraction*. There are three basic relations for abstracting data in heuristic programs:

²For simplicity, we will refer to classification hierarchies throughout this paper, though in practice these structures are not trees, but almost always "tangled" structures with some nodes having multiple parents.

- *definitional abstraction* based on essential, necessary features of a concept ("if the structure is a one-dimensional network, then its shape is a beam");
- *qualitative abstraction*, a form of definition involving quantitative data, usually with respect to some normal or expected value ("if the patient is an adult and white blood count is less than 2500, then the white blood count is low"); and
- *generalization* in a subtype hierarchy ("if the client is a judge, then he is an educated person").

These interpretations are usually made by the program with certainty; belief thresholds and qualifying conditions are chosen so the abstraction is categorical. It is common to refer to this knowledge as being "factual" or "definitional."

2.3. Heuristic classification

In simple classification, data may directly match solution features or may match after being abstracted. In heuristic classification, solutions and solution features may also be matched *heuristically*, by direct, non-hierarchical association with some concept in *another* classification hierarchy. For example, MYCIN does more than identify an unknown organism in terms of visible features of an organism: MYCIN heuristically relates an abstract characterization of the patient to a classification of diseases. We show this *inference structure* schematically, followed by an example (Figure 2-2).

Basic observations about the patient are abstracted to patient categories, which are heuristically linked to diseases and disease categories. While only a subtype link with E.coli infection is shown here, evidence may actually derive from a combination of inferences. Some data might directly match E.coli features (an individual organism shaped like a rod and producing a Gram-negative stain is seen growing in a culture taken from the patient). Descriptions of laboratory cultures (describing location, method of collection, and incubation) can also be related to the classification of diseases.

The important link we have added is a heuristic association between a characterization of the patient ("compromised host") and categories of diseases ("gram-negative infection"). Unlike definitional and hierarchical inferences, this inference makes a great leap. A *heuristic relation* is uncertain, based on assumptions of typicality, and is sometimes just a poorly understood correlation. A heuristic is often empirical, deriving from problem-solving experience; heuristics correspond to the "rules of thumb," often associated with expert systems (Feigenbaum, 1977).

Heuristics of this type reduce search by skipping over intermediate relations (this is why we don't call abstraction relations "heuristics"). These associations are usually uncertain because the intermediate relations may not hold in the specific case. Intermediate relations may be omitted because they are unobservable or poorly understood. In a medical diagnosis program, heuristics typically skip over the causal relations between symptoms and diseases. In Section

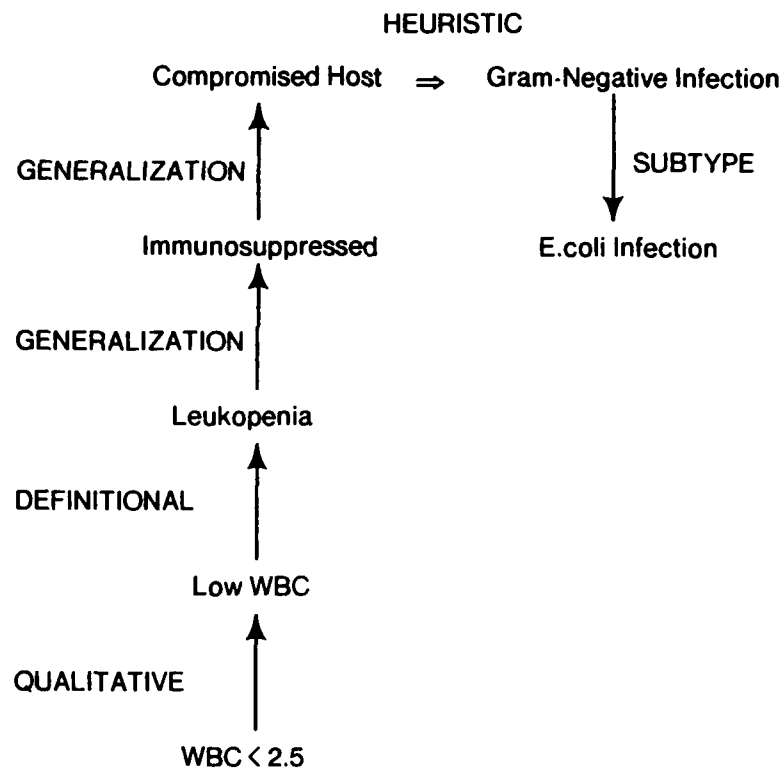
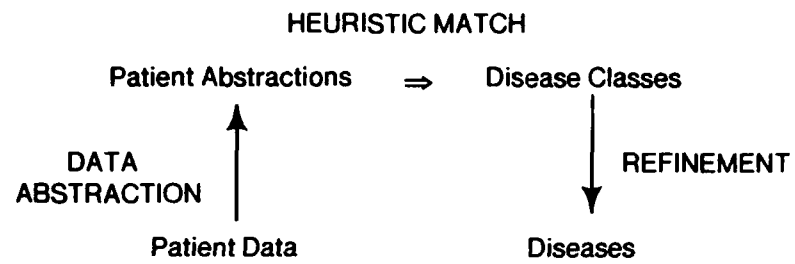


Figure 2-2: Inference structure of MYCIN

4 we will analyze the nature of these implicit relations in some detail.

To summarize, in heuristic classification abstracted data statements are associated with specific problem solutions or features that characterize a solution. This can be shown schematically in simple terms (Figure 2-3).

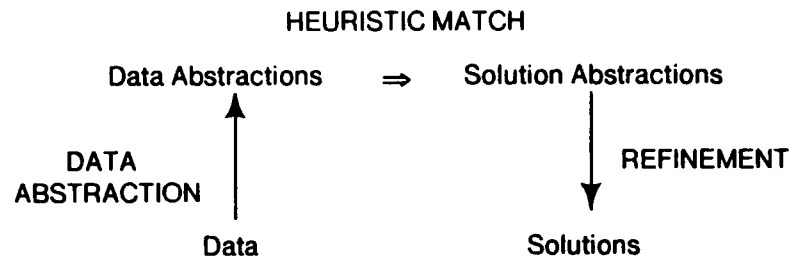


Figure 2-3: Inference structure of heuristic classification

This diagram summarizes how a distinguished set of terms (data, data abstractions, solution abstractions, and solutions) are related systematically by *different* kinds of relations. This is the *structure of inference* in heuristic classification. The direction of inference and the relations "abstraction" and "refinement" are a simplification, indicating a common ordering (generalizing data and refining solutions), as well as a useful way of remembering the classification model. In practice, there are many operators for selecting and ordering inferences, discussed in Section 6.

3. EXAMPLES OF HEURISTIC CLASSIFICATION

Here we schematically describe the architectures of SACON, GRUNDY, and SOPHIE III in terms of heuristic classification. These are brief descriptions, but reveal the value of this kind of analysis by helping us to understand what the programs do. After a statement of the problem, the general inference structure and an example inference path are given, followed by a brief discussion. In looking at these diagrams, note that sequences of classifications can be composed, perhaps involving simple classification at one stage (SACON) or omitting "abstraction" or "refinement" (GRUNDY and SACON).

In the Section 4, we will reconsider these examples, in an attempt to understand the heuristic classification pattern. Our approach will be to pick apart the "inner structure" of concepts and to characterize the kinds of relations that are typically useful for problem solving.

3.1. SACON

Problem: SACON (Bennett, et al., 1978) selects classes of behavior that should be further investigated by a structural-analysis simulation program (Figure 3-1).

Discussion: SACON solves two problems by classification—heuristically analyzing a structure and then using simple classification to select a program. It begins by heuristically selecting a simple numeric model for analyzing a structure (such as an airplane wing). The numeric model, an equation, produces stress and deflection estimates, which the program then qualitatively abstracts as behaviors to study in more detail. These behaviors, with additional information about the material, definitionally characterize different configurations of the MARC simulation program (e.g., the inelastic-fatigue program). There is no refinement because the solutions to the first problem are just a simple set of possible models, and the second problem is only solved to the point of specifying program classes. (In another software configuration system we analyzed, specific program input parameters are inferred in a refinement step.)

3.2. GRUNDY

Problem: GRUNDY (Rich, 1979) is a model of a librarian, selecting books a person might like to read.

Discussion: GRUNDY solves two classification problems heuristically, classifying a reader's personality and then selecting books appropriate to this kind of person (Figure 3-2). While some evidence for people stereotypes is by data abstraction (a JUDGE can be inferred to be an EDUCATED-PERSON), other evidence is heuristic (watching no TV is neither a necessary nor sufficient characteristic of an EDUCATED-PERSON).

Illustrating the power of a knowledge-level analysis, we discover that the people and book classifications are not distinct in the implementation. For example, "fast plots" is a book characteristic, but in the implementation "likes fast plots" is associated with a person stereotype. The relation between a person stereotype and "fast plots" is heuristic and should be distinguished from abstractions of people and books. One objective of the program is to learn better people stereotypes (user models). The classification description of the user modeling problem shows that GRUNDY should also be learning better ways to characterize books, as well as improving its heuristics. If these are not treated separately, learning may be hindered. This example illustrates why a knowledge-level analysis should precede representation.

It is interesting to note that GRUNDY does not attempt to perfect the user model before recommending a book. Rather, refinement of the person stereotype occurs when the reader rejects book suggestions. Analysis of other programs indicates that this multiple-pass process structure is common. For example, the Drilling Advisor makes two passes on the causes of drill sticking, considering general, inexpensive data first, just as medical programs commonly consider the "history and physical" before laboratory data. The high-level, abstract structure of

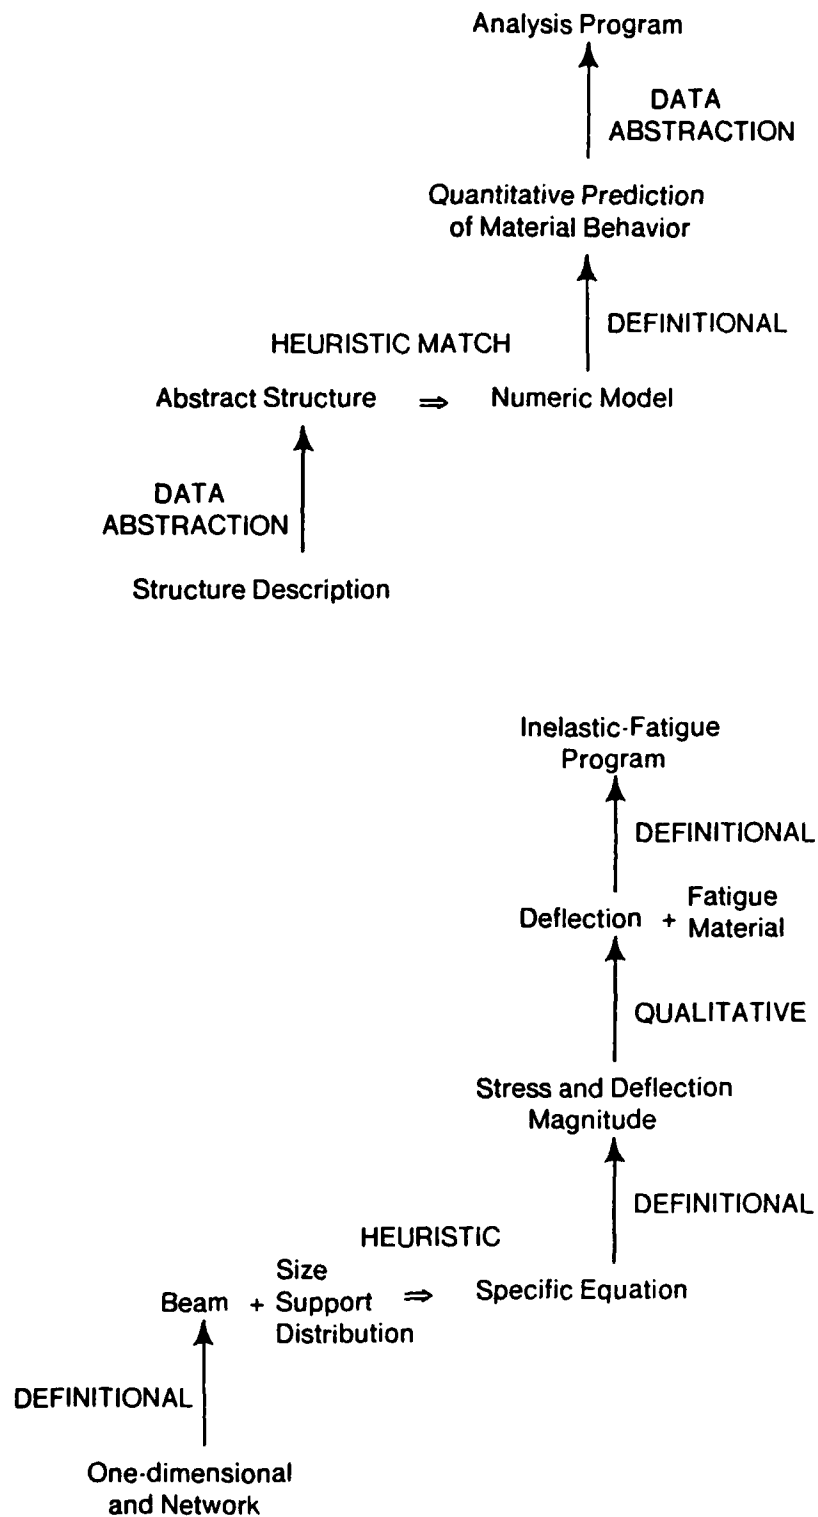


Figure 3-1: Inference structure of SACON

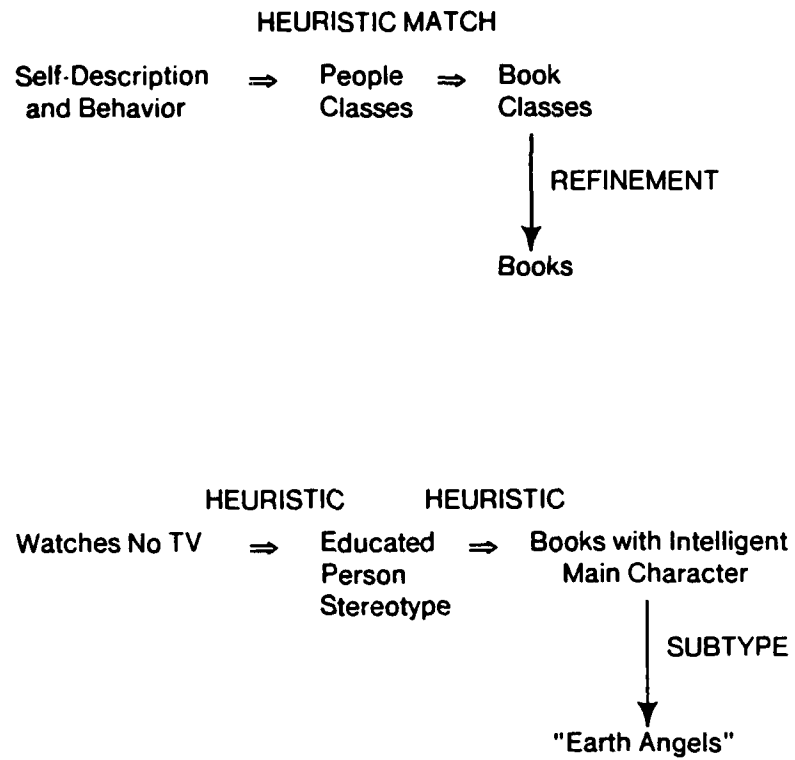


Figure 3-2: Inference structure of GRUNDY

the heuristic classification model makes possible these kinds of descriptions and comparisons.

3.3. SOPHIE III

Problem: SOPHIE III (Brown, et al., 1982) classifies an electronic circuit in terms of the component that is causing faulty behavior (Figure 3-3).

Discussion: SOPHIE's set of pre-enumerated solutions is a lattice of valid and faulty circuit behaviors. In contrast with MYCIN, SOPHIE's solutions are device states and component flaws, not stereotypes of disorders. They are related causally, not by subtype. Data are not only external device behaviors, but include internal component measurements propagated by the causal analysis of the LOCAL program. Nevertheless, the inference structure of abstractions, heuristic relations, and refinement fits the heuristic classification model, demonstrating its generality and usefulness.

4. UNDERSTANDING HEURISTIC CLASSIFICATION

The purpose of this section is to develop a principled account of why the inference structure of heuristic classification takes the characteristic form we have discovered. Our approach is to describe what we have heretofore loosely called "classes," "concepts," or "stereotypes" in a more formal way, using the conceptual graph notation of Sowa (Sowa, 1984). In this formalism, a concept is described by graphs of typed, usually binary relations among other concepts. This kind of analysis has its origins in semantic networks (Quillian, 1968), the conceptual-dependency notation of Schank, et al. (Schank, 1975), the prototype/perspective descriptions of KRL (Bobrow and Winograd, 1979), the classification hierarchies of KL-ONE (Schmolze and Lipkis, 1983), as well as the predicate calculus.

Our discussion has several objectives:

- to relate the knowledge encoded in rule-based systems to structures more commonly associated with "semantic net" and "frame" formalisms,
- to explicate what kinds of knowledge heuristic rules leave out (and thus their advantages for search efficiency and limitations for correctness), and
- to relate the kinds of conceptual relations collectively identified in knowledge representation research (e.g., the relation between an individual and a class) with the pattern of inference that typically occurs during heuristic classification problem solving (yielding the characteristic inverted horseshoe inference structure of Figure 2-3).

One important result of this analysis is a characterization of the "heuristic relation" in terms of primitive relations among concepts (such as preference, accompaniment, and causal

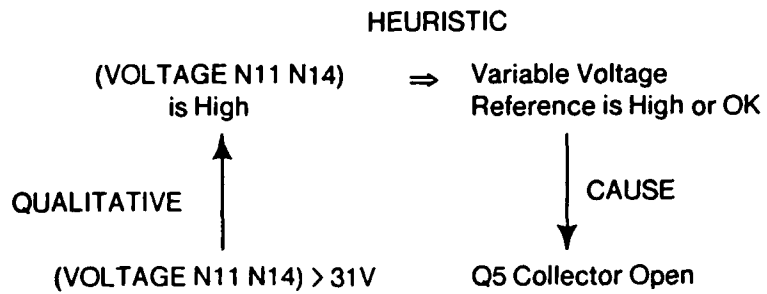
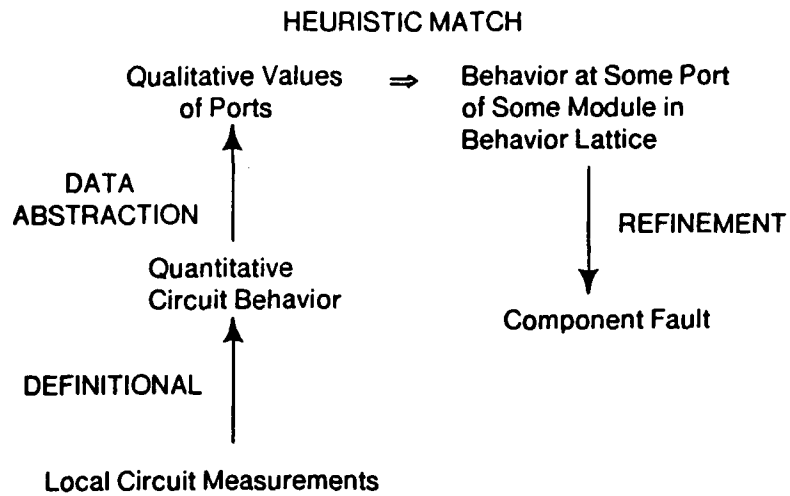


Figure 3-3: Inference structure of SOPHIE

enablement), and its difference from more essential, "definitional" characterizations of concepts. In short, we are trying to systematically characterize *the kind of knowledge that is useful for problem solving*, which relates to our larger aim of devising useful languages for encoding knowledge in expert systems.

4.1. Schemas vs. definitions

In the case of matching features of organisms (MYCIN) or programs (SACON), features are essential (necessary), identifying characteristics of the object, event, or process. This corresponds to the Aristotelian notion of concept definition in terms of necessary properties.³ In contrast, features may be only "incidental," corresponding to typical manifestations or behaviors. For example, E.coli is normally found in certain parts of the body, an incidental property. It is common to refer to the combination of incidental and defining associations as a "schema" for the concept.⁴ Inferences made using incidental associations of a schema are inherently uncertain. For example, we might infer that a particular person, because he is educated, likes to read books, but this might not be true. In contrast, an educated person must, by definition, have learned a great deal about something (though maybe not a formal academic topic).

The nature of schemas and their representation has been studied extensively in AI. As stated in the introduction (Section 1), our purpose here is to exploit this research to understand the knowledge contained in rules. We are not advocating one representation over another; rather we just want to find some way of writing down knowledge so that we can detect and express patterns. We use the conceptual graph notation of Sowa because it is simple and it makes basic distinctions that we find to be useful:

- A schema is made up of *coherent statements* mentioning a given concept, not a list of isolated, independent features. (A statement is a complete sentence.)
- A schema for a given concept contains *relations to other concepts*, not just "attributes and values" or "slots and values."
- A concept is typically described from different points of view by *a set of schemata* (called a "schematic cluster"), not a single "frame."

³ (Sowa, 1984) provides a good overview of these well-known philosophical distinctions. See also (Palmer, 1978) and (Cohen and Murphy, 1984).

⁴ Here we use the word "schema" as a *kind of knowledge*, not a construct of a particular programming language or notation. See (Hayes, 1979) for further discussion of this distinction.

- The totality of what people know about a concept usually extends well beyond the schemas that are pragmatically encoded in programs for solving limited problems.

Finally, we adopt Sowa's definition of a *prototype* as a "typical individual," a specialization of a concept schema to indicate typical values and characteristics, where ranges or sets are described for the class as a whole. Whether a program uses prototype or schema descriptions of its solutions is not important to our discussion, and many may combine them, including "normal" values, as well as a spectrum of expectations.

4.2. Alternative encodings of schemas

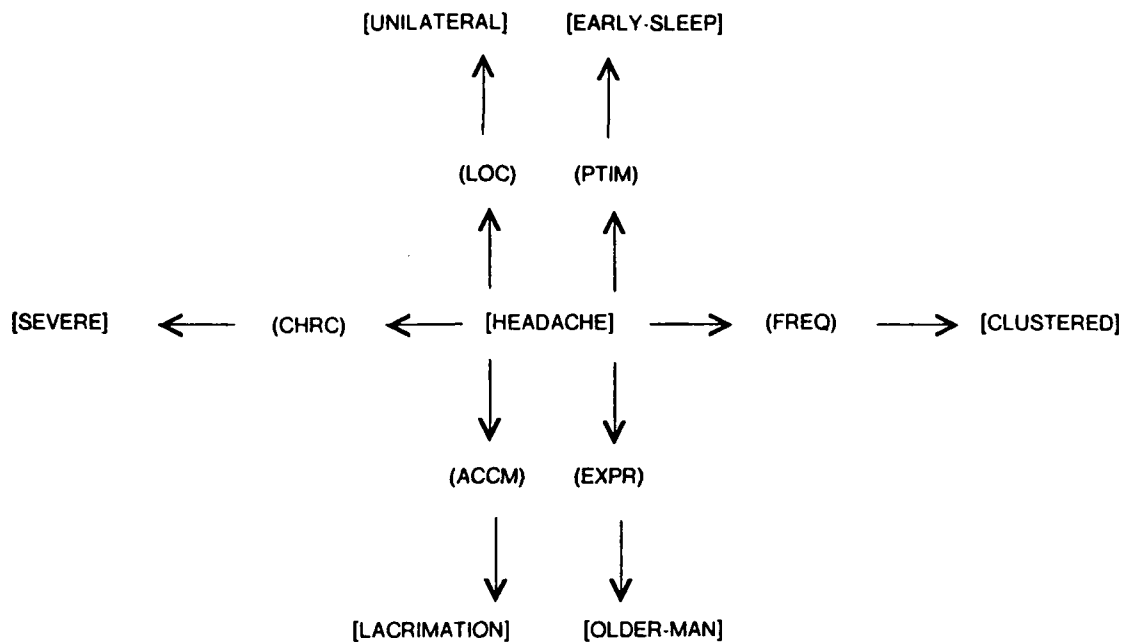
To develop the above points in some detail, we will consider a conceptual graph description and how it relates to typical rule-based encodings. Figure 4-1 shows how knowledge about the concept "cluster headache" is described using the conceptual graph notation.⁵

Concepts appear in brackets; relations are in parentheses. Concepts are also related by a type hierarchy, e.g., a HEADACHE is a kind of PROCESS, an OLDER-MAN is a kind of MAN. Relations are constrained to link concepts of particular types, e.g., PTIM, a point in time, links a PROCESS to a TIME. For convenience, we can also use Sowa's linear notation for conceptual graphs. Thus, OLDER-MAN can be described as a specialization of MAN, "a man with characteristic old." CLUSTERED is "an event occurring daily for a week." EARLY-SLEEP is "a few hours after the state of sleep."

We make no claim that a representation of this kind is complete, computationally tractable, or even unambiguous. For our purposes here, it is simply a *notation* with the advantage over English prose of systematically revealing how what we know about a concept can be (at least partially) described in terms of its relations to concepts of other types.

For contrast, consider how this same knowledge might be encoded in a notation based upon objects, attributes, and values, as in MYCIN. Here, the object would be the PATIENT, and typical attributes would be HEADACHE-ONSET (with possible values EARLY-MORNING, EARLY-SLEEP, LATE-AFTERNOON) and DISORDER (with possible values CLUSTER-HEADACHE, INFECTION, etc.). A typical rule might be, "If the patient has headache onset during sleep, then the disorder of the patient is cluster headache." The features of a cluster headache might be combined in a single rule. Generally, since none of the features are logically necessary, they are considered in separate rules, with certainty factors denoting how strongly the symptom (or predisposition, in the case of age) is correlated with the disease. A

⁵One English translation would be: "A cluster headache is a headache that occurs with a frequency in clusters, experienced by an older man, accompanied by lacrimation, with characteristic severe, of location unilateral, occurring at a point in time of early sleep."



[EARLY-SLEEP] is
 [TIME: [STATE: [SLEEP]] -> (AFTER) -> [TIME-PERIOD: @few-hrs]]

[CLUSTERED] is
 [DAILY] <- (FREQ) <- [EVENT] -> (DURATION) -> [TIME-PERIOD: @1week]

[OLDER-MAN] is
 [MAN] -> (CHRC) -> [OLD]

Figure 4-1: Schema describing the concept CLUSTER-HEADACHE and some related concepts

primitive "frame" representation, as in INTERNIST (Pople, 1982), is similar, with a list of attributes for each disorder, but each attribute is an "atomic" unit that bundles together what is broken into object, attribute, and value in MYCIN, e.g., "HEADACHE-ONSET-OCCURS-EARLY-SLEEP."

The idea of relating a concept (such as CLUSTER-HEADACHE) to a set of attributes or descriptors, is common in AI programs. However, a relational analysis reveals marked differences in what an attribute might be:

- An attribute is an atomic proposition. In INTERNIST, an attribute is a string that is only related to diseases or other strings, e.g., HEADACHE-ONSET-EARLY-SLEEP-EXPERIENCED-BY-PATIENT.

- **An attribute is a relation characterizing some class of objects.** In MYCIN, an attribute is associated with an instance of an object (a particular patient, culture, organism, or drug).
 - **An attribute is a unary relation.** A MYCIN attribute with the values "yes or no" corresponds to a unary relation, ($\langle \text{attribute} \rangle \langle \text{object} \rangle$), e.g., (HEADACHE-ONSET-EARLY-SLEEP PATIENT), "headache onset during early sleep is experienced by the patient."
 - **An attribute is a binary relation.** A MYCIN attribute with *values* corresponds to a binary relation, ($\langle \text{attribute} \rangle \langle \text{object} \rangle \langle \text{value} \rangle$), e.g., (HEADACHE-ONSET PATIENT EARLY-SLEEP), "headache onset experienced by the patient is during early sleep."
- **An attribute is a relation among classes.** Each class is a concept. Taking the same example, there are two more primitive relations, ONSET and EXPERIENCER, yielding the propositions: (ONSET HEADACHE EARLY-SLEEP), "the onset of the headache is during early sleep", and (EXPERIENCER HEADACHE PATIENT), "the experiencer of the headache is the patient." More concisely, [EARLY-SLEEP] \leftarrow (ONSET) \leftarrow [HEADACHE] \rightarrow (EXPR) \rightarrow [PATIENT]. These relations and concepts can be further broken down, as shown in Figure 4-1.

The conceptual graph notation encourages clear thinking by forcing us to unbundle domain terminology into defined or schematically described terms and a constrained vocabulary of relations (restricted in the types of concepts each can link). Rather than saying that "an object has attributes," we can be more specific about the relations among entities, describing abstract concepts like "headache" and "cluster" in the same notation we use to describe concrete objects like patients and organisms. In particular, notice that headache onset is a characterization of a headache, not of a person, contrary to the MYCIN statement that "headache onset is an attribute of person." Similarly, the relation between a patient and a disorder is different from the relation between a patient and his age.⁶

Breaking apart "parameters" into concepts and relations has the additional benefit of allowing them to be easily related, through their schema descriptions. For example, it is clear that HEADACHE-ONSET and HEADACHE-SEVERITY both characterize HEADACHE, allowing

⁶The importance of defining relations has been discovered repetitively in AI. Wood's analysis of semantic networks (Woods, 1975) is an early, well-known example. The issue of restricting and defining relations was particularly important in the development of OWL (Martin, 1979). Researchers using rule-based languages, like MYCIN's, felt curiously immune from these issues, not realizing that their "attributes" were making similar confusions.

us to write a simple, general inference rule for deciding about relevancy: "If a process type being characterized (e.g., HEADACHE) is unavailable or not relevant, then its characterization (e.g., HEADACHE-ONSET) is not relevant." As another example, consider a discrimination inference strategy that compares disorder processes on the basis of their descriptions as *events*. Knowing what relations are comparable (e.g., location and frequency), the inference procedure can automatically gather relevant data, look up the schema descriptions, and make comparisons to establish the best match. To summarize, the rules in a program like MYCIN are implicitly making statements about schemas. This becomes clear when we separate conceptual links from rules of inference, as in NEOMYCIN.

4.3. Relating heuristics to conceptual graphs

Given all of the structural and functional statements we might make about a concept, describing processes and interactions in detail, some statements will be more useful than others for solving problems. Rather than thinking of schemas as inert, static descriptions, we are interested in how they link concepts to solve problems. The description of CLUSTERED-HEADACHE given in Figure 4-1 includes the knowledge that one typically finds in a diagnostic program. To understand heuristics in these terms, consider first that some relations appear to be less "incidental" than others. The time of occurrence of the headache, location, frequency, and characterizing features are all closely bound to what a cluster headache is. They are not necessary, but they together distinguish CLUSTER-HEADACHE from other types. That is, these relations *discriminate* this headache from other types of headache.

On the other hand, accompaniment by lacrimation (tearing of the eyes) and the tendency for such headaches to be experienced by older men are correlations with other concepts.⁷ Here, in particular, we see the link between different kinds of entities: a DISORDER-PROCESS and a PERSON. This is the link we have identified as a heuristic—a direct, non-hierarchical association between concepts of different types. Observe that *why* an older man experiences cluster headaches is left out. Given a model of the world that says that all phenomena are caused, we can say that each of the links with HEADACHE could be explained causally. Whether the explanation has been left out or is not known cannot be determined by examining the conceptual graph, a critical point we will return to later.

When heuristics are stated as rules in a program like MYCIN, even known relational and definitional details are often omitted. This often means that intermediate concepts are omitted as well. We say "X suggests Y, or "X makes you think of Y." Unless the connection is an

⁷What discriminates is relative. If kinds of headache tended to be associated with different ages of people, then this might be a CLUSTER-ELDERLY-HEADACHE and we would consider the age of the experiencer to be a discriminating characteristic.

unexplained correlation, such a statement can be expanded to a full sentence that is part of the schema description of X and/or Y. Thus, the geologist's rule "goldfields flowers --> serpentine rock" might be restated as, "Serpentine rock has nutrients that enable goldfields to grow well." Figure 4-2 shows the conceptual graph notation of this statement (with "enable" shown by the relation "instrument" linking an entity, nutrients, to an act, growing).




Figure 4-2: A heuristic rule expanded as a conceptual graph

The concepts of nutrients and growing are omitted from the rule notation, just as the causal details that explain the growth process are skipped over in the conceptual graph notation. The rule indicates what you must observe (goldfields flowers growing) and what you can assert (serpentine rock is near the surface). It captures knowledge not as mere static descriptions, but as efficient, useful connections for problem solving. Moreover, the example makes clear the essential characteristic of a heuristic inference—a non-hierarchical and non-definitional connection between concepts of distinct classes.

Heuristics are selected statements that are useful for inference, particularly how one class choice *constrains* another. Consider the goldfields example. Is the conceptual graph shown in Figure 4-2 a schema for serpentine, goldfields, nutrient, or all three? First, knowledge is something somebody knows; whether goldfields is associated with nutrients will vary from person to person. (And for at least a short time, readers of this paper will think of goldfields when the word "nutrient" is mentioned.) Second, the real issue is how knowledge is *practically indexed*. The associations a problem solver forms and the directionality of these associations will depend on the kinds of situations he is called upon to interpret, and what is given and what is derived. Thus, it seems plausible that a geologist in the field would see goldfields (data) and think about serpentine rock (solution). Conversely, his task might commonly be to find outcroppings of serpentine rock; he would work backwards to think of observables that he might look for (data) that would indicate the presence of serpentine. Indeed, he might have many associations with flowers and rocks, and even many general rules for how to infer rocks (e.g., based on other plants, drainage properties of the land, slope). Figure 4-3 shows one possible inference path.

In summary, a heuristic association is a connection that relates data that is commonly available to the kinds of interpretations the problem solver is trying to derive. For a

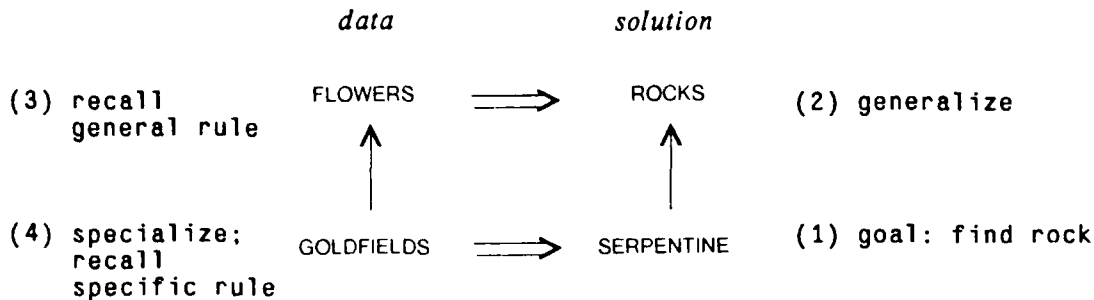


Figure 4-3: Using a general rule to work backwards from a solution

physician starting with characteristics of a person, the patient, connections to diseases will be useful. It must be possible to relate new situations to previous interpretations and this is what the abstraction process in classification is all about (recall the quotation from Bruner in Section 1). The specific person becomes an "old man" and particular disorders come to mind.

Problems tend to start with objects in the real world, so it makes sense that practical problem-solving knowledge would allow problems to be restated in terms of stereotypical objects: kinds of people, kinds of patients, kinds of stressed structures, kinds of malfunctioning devices, etc. Based on our analysis of expert systems, links from these data concepts to solution concepts come in different flavors:

- agent or experiencer (e.g., people predisposed to diseases)
- cause, co-presence, or correlation (e.g., symptoms related to faults)
- preference or advantage (e.g., people related to books)
- physical model (e.g., abstract structures related to numeric models)

These relations don't characterize a solution in terms of "immediate properties"—they are not definitional or type discriminating. Rather, they capture *incidental associations between a solution and available data, usually concrete concepts*. (Other kinds of links may be possible; these are the ones we have discovered so far.)

The essential characteristic of a heuristic is that it reduces search. A heuristic rule reduces a conceptual graph to a single relation between two concepts. Through this, heuristic rules reduce search in several ways:

1. Of all possible schemas that might describe a concept, heuristic connections are those that constrain a categorization on the basis of available data (e.g., the strength of SERPENTINE rock may be irrelevant for inferring the presence of hidden deposits).

2. A heuristic eliminates consideration of intermediate (and often invariant) relations between the concepts it mentions, associating salient classes directly (e.g., the goldfields rule omits the concept NUTRIENT).

While not having to think about intermediate connections is advantageous, this sets up a basic conflict for the problem solver—his inferential leaps may be wrong. Another way of saying that the problem solver skips over things is that there are *unarticulated assumptions* on which the interpretation rests. We will consider this further in the section on inference strategies (Section 6).

4.4. Relating inference structure to conceptual graphs

In the inference-structure diagrams (such as Figure 3-2) nodes stand for propositions (e.g., "the reader is an educated person"). The diagrams relate propositions on the basis of how they can be inferred from one another: type, definition, and heuristic. So far in this section we have broken apart these atomic propositions to distinguish a heuristic link from essential and direct characterizing relations in a schema; and we have argued how direct, accidental connections between concepts, which leave out intermediate relations, are valuable for reducing search.

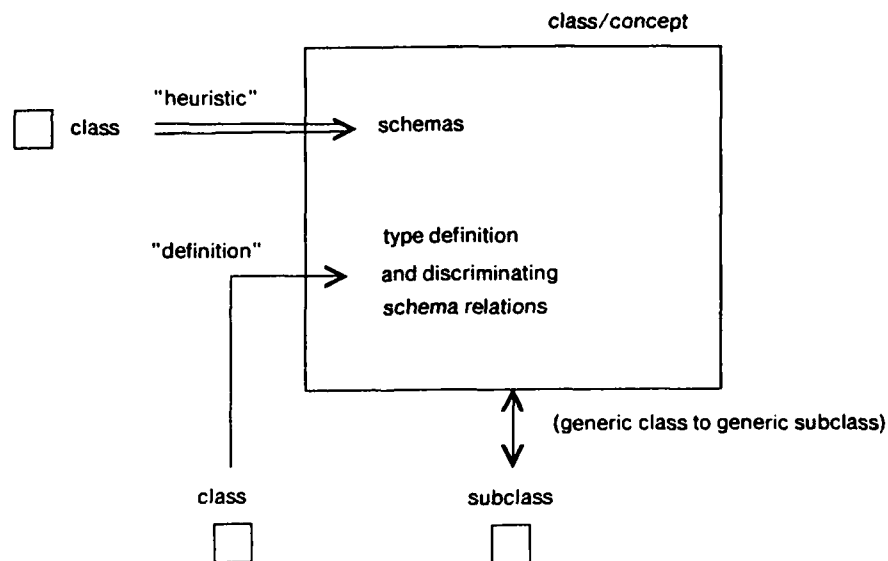


Figure 4-4: Conceptual relations used in heuristic classification

Here we return to the higher-level, inference-structure diagrams and include the details of the kinds of links that are possible. In Figure 4-4 each kind of inference relation between

concepts is shown as a line. Classes can be connected to one another by any of these three kinds of inference relations. We make a distinction between heuristics (direct, non-hierarchical, class-class relations, such as the link between goldfields and serpentine rock) and definitions (including necessary and discriminating relations, plus qualitative abstraction (see Section 2.2)). Definitional and subtype links are shown vertically, to conform to our intuitive idea of generalization of data to higher categories, what we have called *data abstraction*.

It is important to remember that the "definitional" links are often non-essential, "soft" descriptions. The "definition" of leukopenia as white blood count less than normal is a good example. "Normal" depends on everything else happening to the patient, so inferring this condition always involves making some assumptions.

Note also that this is a diagram of static, structural relations. In actual problem solving other links will be required to form a *case-specific model*, indicating propositions the problem solver believes to be true and support for them. In particular, *surrogates* (Sowa, 1984) (also called *individuals* (Brachman, 1977), such the MYCIN "context" ORGANISM-1) will stand for unknown objects or processes in the world that are identified by associating them with a class in a type hierarchy.⁸

Now we are ready to put this together to understand the pattern behind the inference structure of heuristic classification. Given that a sequence of heuristic classifications, as in GRUNDY, is possible, indeed common, we start with the simplest case by assuming that data classes are not inferred heuristically. Instead, data are supplied directly or inferred by definition. When solution classes are inferred by definition, we have a case of *simple classification* (Section 2.1), for example, when an organism is actually seen growing in a laboratory culture (like a smoking gun). In order to describe an idealized form of heuristic

⁸It is not often realized that each MYCIN "context" has a *distinguished attribute* called its "name" that corresponds to the link between the surrogate (entity to be classified) and a classification hierarchy. The pattern was only evident to system designers to the extent that they realized that each "context" type has some identifying attribute that allows it to be translated. For example, after identifying an organism, the program says "the E.coli" rather than ORGANISM-1 (or whatever its number was), referring to the object/context hierarchy if there are more than one, "the E.coli from the blood culture of 3/14/77." Thus, we have the identity of the organism, name of the infection, site of the culture, etc. Corresponding to each of these identifying attributes is a hierarchy of "values" with *static properties*. Thus, there are tables of organisms, infections, culture sites, etc. It is in such a table that MYCIN stores the information that E.coli is a gram-negative rod. A single, general rule uses the table to identify the unknown organism. These tables are also called "grids"; we were unaware at the time (1974-1977) that we were recording the same kind of information other AI programmers were storing in "frame hierarchies." The pattern was partially obscured by our use of special-case rules, for example, to allow for incorrect data, making the grids appear to be a convenient computational short-hand for collapsing similar rules, rather than a notation for describing classes.

classification, we leave out definitional inference of solutions. Finally, inference has the general form that problem descriptions must be abstracted (proceeding from subclass to class) and partial solutions must be refined (proceeding from class to subclass).

If we thus specialize the right side of the inference diagram in Figure 4-4 to a data class and a solution class and glue them together, we get a refined version of the original inverted horseshoe (Figure 2-3). Figure 4-5 shows how data and solution classes are typically inferred from one another in the simplest case of heuristic classification. This diagram should be contrasted with all of the possible networks we could construct, linking concepts by the three most general relations (subtype, definitional, incidental). For example, all links might have been definitional, all concepts subsumed by a single class, or data only incidentally related to other concepts. Furthermore, considering knowledge apart from how it is used, we might imagine complex networks of concepts, intricately related, as suggested by Figure 4-1. Instead, we find that diverse classification structures are often linked directly, omitting relational details. Clearly independent of programming language, this pattern is very likely an essential aspect of practical, experiential models of the world.

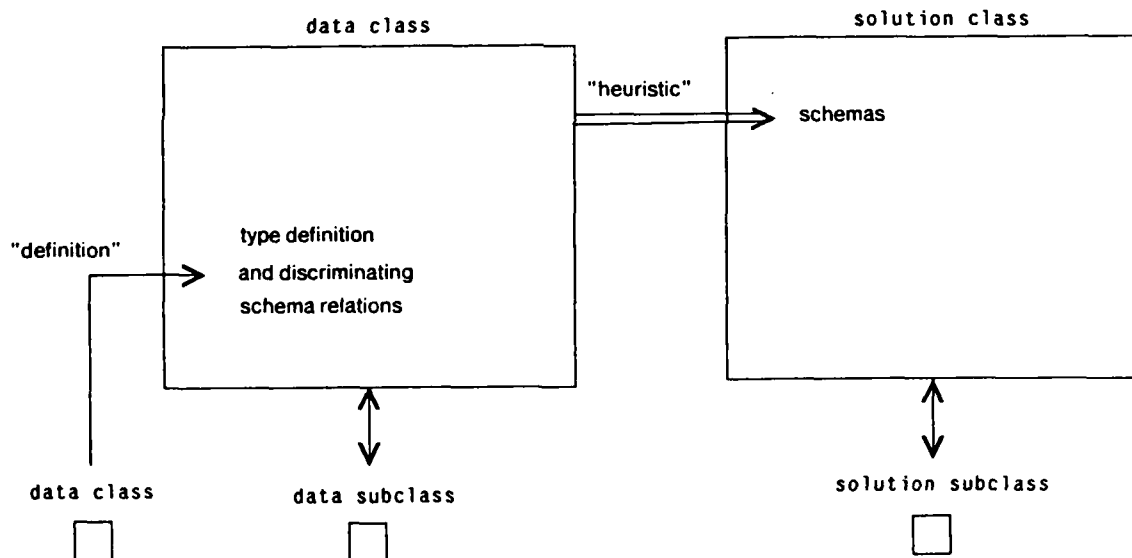


Figure 4-5: Typical conceptual relations in simplest form of heuristic classification

4.5. Pragmatics of defining concepts

In the course of writing and analyzing heuristic programs, we have been struck by the difficulty of defining terms. What is a "compromised host?" How is it different from "immunosuppression"? Is an alcoholic immunosuppressed? We do not simply write down descriptions of what we know. The very process of formalizing terms and relations changes what we know, and itself brings about concept formation.

In many respects, the apparent unprincipled nature of MYCIN is a good reflection of the raw state of how experts talk. Two problems we encountered illustrate the difficulty of proceeding without a formal conceptual structure, and thus, reflect the unprincipled state of what experts know about their own reasoning:

- Twice we completely reworked the hierarchical relations among immunosuppression and compromised host conditions. There clearly is no agreed-upon network that we can simply write down. People do not know schema hierarchies in the same sense that they know phone numbers. A given version is believed to be better because it makes finer distinctions, so it leads to better problem solving.
- The concepts of "significant organism" and "contaminant" were sometimes confused in MYCIN. An organism is significant if there is evidence that it is associated with a disease. A contaminant is an organism growing on a culture that was introduced because of dirty instruments or was picked up from some body site where it normally grows (e.g., a blood culture may be contaminated by skin organisms). Thus, evidence against contamination supports the belief that the discovered organism is significant. However, a rule writer would tend to write "significant" rather than "not contaminant," even though this was the intended, intermediate interpretation. There may be a tendency to directly form a general, positive categorization, rather than to make an association to an intermediate, ruled-out category.

To a first approximation, it appears that what we "really" know is what we can conclude given other information. That is, we start with just implication ($P \rightarrow Q$), then go back to abstract concepts into types and understand relations among them. For example, we start by knowing that "WBC < 2500 \rightarrow LEUKOPENIA." To make this principled, we break it into the following pieces:

1. "Leukopenia" means that the count of leukocytes is impoverished:
 $[LEUKOPENIA] = [LEUKOCYTES] \rightarrow (CHRC) \rightarrow [CURRENT-COUNT] \rightarrow$
 $(CHRC) \rightarrow [IMPOVERISHED]$

2. "Impoverished" means that the current measure is much less than normal:

[IMPOVERISHED: x] =

[CURRENT-MEASURE: x] -> (<<) -> [NORMAL-MEASURE: x]

3. The (normal/current) count is a kind of measure:

[COUNT] < [MEASURE]

4. A fact, the normal leukocyte count in an adult is 7000:

[LEUKOCYTES] -> (CHRC) ->

[NORMAL-COUNT] -> (MEAS) ->

[MEASURE: 7000 /mm3].

With the proper interpreter (and perhaps some additional definitions and relations), we could instantiate and compose these expressions to get the effect of the original rule. This is the pattern we follow in knowledge engineering, constantly decomposing terms into general types and relations to make explicit the rationale behind implications.

Perhaps one of the most perplexing difficulties we encounter is distinguishing between subtype and cause, and between state and process. Part of the problem is that cause and effect are not always distinguished by our experts. For example, a physician might speak of a brain-tumor as a kind of brain-mass-lesion. It is certainly a kind of brain mass, but it causes a lesion (cut); it is not a kind of lesion. Thus, the concept bundles cause with effect and location: a *lesion* in the *brain* caused by a *mass* of some kind is a brain-mass-lesion (Figure 4-6).

[MASS] -> (CAUS) -> [LESION] -> (LOC) -> [BRAIN]

Figure 4-6: Conceptual graph of the term "brain-mass-lesion"

Similarly, we draw causal nets linking abnormal states, saying that brain-hematoma (mass of blood in the brain) is caused by brain-hemorrhage (bleeding). To understand what is happening, we profit by labeling brain-hematoma as a *substance* (a kind of brain-mass) and brain-hemorrhage as a *process* that affects or produces the substance. Yet when we began, we thought of brain-hemorrhage as if it were equivalent to the escaping blood.

It is striking that we can learn concepts and how to relate them to solve problems, without understanding the links in a principled way. If you know that $WBC < 2500$ is leukopenia, a form of immunosuppression, which is a form of compromised host, causing E.coli infection, you are on your way to being a clinician. As novices, we push tokens around in the same non-comprehending way as MYCIN.

Once we start asking questions, we have difficulty figuring out how concepts are related. If immunosuppression is the state of being unable to fight infection by mechanisms, then does impoverished white cells *cause* this state? Or is it *caused by* this state (something else affected the immunosystem, reducing the WBC as a side-effect)? (Worse yet, we may say it is an "indicator," completely missing the fact that we are talking about causality.) Perhaps it is one way in which the immunosystem can be diminished, so it is *a kind of* immunosuppression. It is difficult to write down a principled network because we don't know the relations, and we don't know them because we don't know what the concepts mean—we don't understand the processes involved. Yet, we might know enough to relate data classes to therapy classes and save the patient's life!

A conceptual graph or logic analysis suggests that the relations among concepts are *relatively* few in number and fixed in meaning, compared to the number and complexity of concepts. The meaning of concepts depends on what we ascribe to the links that join them. Thus, in practice we jockey around concepts to get a well-formed network. Complicating this is our tendency to use terms that bundle cause with effect and to relate substances directly, leaving out intermediate processes. At first, novices might be like today's expert programs. A concept is just a token or label, associated with knowledge of how to infer truth and how to use information (what to do if it is true and how to infer other truths from it). Unless the token is defined by something akin to a conceptual graph, it is difficult to say that the novice or program understands what it means. But in the world of action, what matters more than the functional, pragmatic knowledge of knowing what to do?

Where does this leave us? One conclusion is that "principled networks" are impossible. Except for mathematics, science, economics, and similar domains, concepts do not have formal definitions. While heuristic programs sometimes reason with concrete, well-defined classifications such as the programs in SACON and the fault network in SOPHIE, they more often use experiential schemas, the knowledge we say distinguishes the expert from the novice. In the worst case, these experiential concepts are vague and incompletely understood, such as the diseases in MYCIN. In general, there are underlying (unarticulated or unexamined) assumptions in every schema description. Thus, the first conclusion is that for concepts in nonformal domains this background and context cannot in principle be made explicit (Flores and Winograd, 1985). That is, our conceptual knowledge is inseparable from our as yet ungeneralized memory of experiences.

An alternative point of view is that, regardless of ultimate limitations, it is obvious that expert systems will be valuable for replacing the expert on routine tasks, aiding him on difficult tasks, and generally transforming how we write down and teach knowledge. Much more can be done in terms of memory representation, learning from experience, and combining principled models with situation/action, pragmatic rules. Specifically, the problem of *knowledge transformation* could become a focus for expert systems research, including

compilation for efficiency, derivation of procedures for enhancing explanation (Swartout, 1981), and re-representation for detecting and explaining patterns, thus aiding scientific theory formation. Studying and refining actual knowledge bases, as exemplified by this section, is our chief methodology for improving our representations and inference procedures. Indeed, from the perspective of knowledge transformation, it is ironic to surmise that we might one day decide that the "superficial" representation of EMYCIN rules is a fine executable language, and something like it will become the target for our knowledge compilers.

5. ANALYSIS OF PROBLEM TYPES IN TERMS OF SYSTEMS

The heuristic classification model gives us two new angles for comparing problem-solving methods and kinds of problems. First, it suggests that we characterize programs by whether solutions are selected or constructed. This leads us to the second perspective, that different "kinds of things" might be selected or constructed (diagnoses, user models, etc.). In this section we will adopt a single point of view, namely that a solution is most generally a set of beliefs describing what is true *about a system* or a set of actions (operations) that will physically *transform a system* to a desired description. We will study variations of system description and transformation problems, leading to a hierarchy of *kinds of problems* that an expert might solve.

5.1. What gets selected?

This foray into systems analysis begins very simply with the observation that all classification problem solving involves *selection* of a solution. We can characterize kinds of problems by *what is being selected*:

- *diagnosis*: solutions are faulty components (SOPHIE) or processes affecting the device (MYCIN);
- *user model*: solutions are people stereotypes in terms of their goals and beliefs (first phase of GRUNDY);
- *catalog selection*: solutions are products, services, or activities, e.g., books, personal computers, careers, travel tours, wines, investments (second phase of GRUNDY);
- *model-based analysis*: solutions are numeric models (first phase of SACON);
- *skeletal planning*: solutions are plans, such as packaged sequences of programs and parameters for running them (second phase of SACON, also first phase of experiment planning in MOLGEN (Friedland, 1979)).

Attempts to make knowledge engineering a systematic discipline often begin with a listing of kinds of problems. This kind of analysis is always prone to category errors. For example, a naive list of "problems" might list "design," "constraint satisfaction," and "model-based reasoning," combining a kind of problem, an inference method, and a kind of knowledge. For example, one might solve a VLSI chip design problem using constraint satisfaction to reason

about models of circuit components. It is important to adopt a single perspective when making a list of this kind.

In particular, we must not confuse what gets selected—what constitutes a solution—with the method for computing the solution. A common misconception is that there is a kind of problem called a "classification problem," opposing, for example, classification problems with design problems (for example, see (Sowa, 1984)). Indeed, some tasks, such as identifying bacteria from culture information, are inherently solved by *simple classification*. However, heuristic classification as defined here is *a description of how a particular problem is solved by a particular problem solver*. If the problem solver has a priori knowledge of solutions and can relate them to the problem description by data abstraction, heuristic association, and refinement, then the problem can be solved by classification. For example, if it were practical to enumerate all of the computer configurations R1 might select, or if the solutions were restricted to a predetermined, explicit set of designs, the program could be reconfigured to solve its problem by classification. The method of solving a configuration problem is not inherent in the task itself.

With this distinction between problem and computational method in mind, we turn our attention to a systematic study of problem types. Can we form an explicit taxonomy that includes the kinds of applications we might typically encounter?

5.2. Background: Problem categories

One approach might be to focus on objects and what can be done to them. We can design them, diagnose them, use them in a plan to accomplish some function, etc. This seems like one way to consistently describe kinds of problems. Surely everything in the world involves objects.

However, in attempting to derive such a uniform framework, the concept of "object" becomes a bit elusive. For example, the analysis of problem types in Building Expert Systems (hereafter BES, (Hayes-Roth, et al., 1983), see Table 5-1) indirectly refers to a program as an object. Isn't it really a process? Are procedures objects or processes? It's a matter of perspective. Projects and audit plans can be thought of as both objects and processes. Is a manufacturing assembly line an object or a process? The idea of a "system" appears to work better than the more common focus on objects and processes.

By organizing descriptions of problems around the concept of a system, we can improve upon the distinctions made in BES. As an example of the difficulties, consider that a situation description is a description of a system. Sensor data are observables. But what is the difference between INTERPRETATION (inferring system behavior from observables) and DIAGNOSIS (inferring system malfunctions from observables)? Diagnosis, so defined, includes interpretation. The list appears to deliberately have this progressive design behind it, as is

INTERPRETATION	Inferring situation descriptions from sensor data
PREDICTION	Inferring likely consequences of given situation
DIAGNOSIS	Inferring system malfunctions from observables
DESIGN	Configuring objects under constraints
PLANNING	Designing actions
MONITORING	Comparing observations to plan vulnerabilities
DEBUGGING	Prescribing remedies for malfunctions
REPAIR	Executing a plan to administer a prescribed remedy
INSTRUCTION	Diagnosis, debugging, and repairing student behavior
CONTROL	Interpreting, predicting, repairing, and monitoring system behaviors.

Table 5-1: Generic categories of knowledge engineering applications.

From

(Hayes-Roth, et al., 1983) Table 1.1, page 14

particularly clear from the last two entries, which are composites of earlier "applications." In fact, this idea of multiple "applications" to something (student behavior, system behavior) suggests that a simplification might be found by adopting more uniform terminology. As a second example, consider that the text of BES says that automatic programming is an example of a problem involving planning. How is that different from configuration under constraints (i.e., design)? Is automatic programming a planning problem or a design problem? We also talk about experiment design and experiment planning. Are the two words interchangeable?

We can get clarity by turning things around, thinking about systems and what can be done to and with them.

5.3. A system-oriented approach

We start by informally defining a *system* to be a complex of interacting objects that have some process (I/O) behavior. The following are examples of systems:

- a stereo system
- a VLSI chip
- an organ system in the human body
- a computer program
- a molecule
- a university
- an experimental procedure

Webster's defines a system to be "a set or arrangement of things so related or connected as to form a unity or organic whole." The parts taken together have some structure. It is useful to think of the unity of the system in terms of how it behaves. Behavior might be characterized simply in terms of inputs and outputs.

Figures 5-1 and 5-2 summarize hierarchically what we can do to or with a system, revising the BES table. We group operations in terms of those that *construct* a system and those that

interpret a system, corresponding to what is generally called *synthesis* and *analysis*. Common synonyms appear in parentheses below the generic operations. In what follows, our new terms appear in upper case.

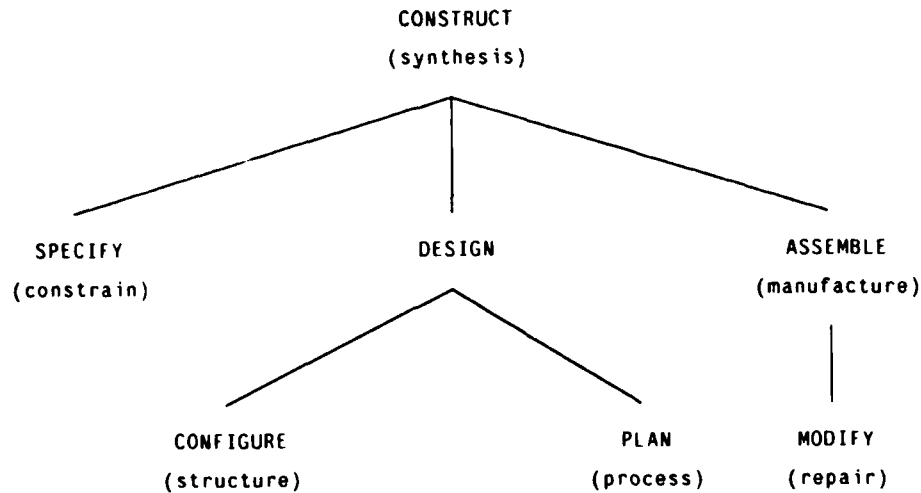


Figure 5-1: Generic operations for synthesizing a system

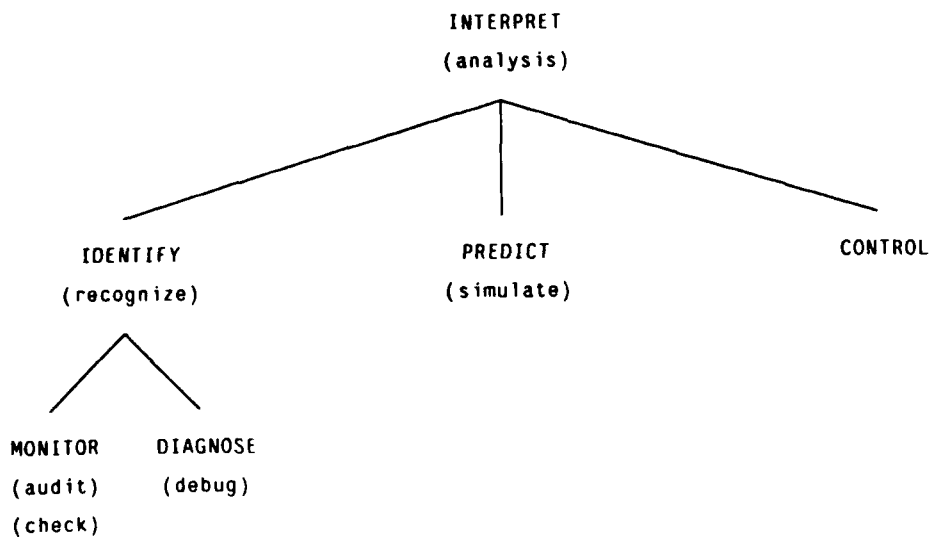


Figure 5-2: Generic operations for analyzing a system

INTERPRET operations concern a working system in some environment. In particular, IDENTIFY is different from DESIGN in that it requires taking I/O behavior and mapping it

onto a system. If the system has not been described before, then this is equivalent to (perhaps only partial) design from I/O pairs. PREDICT is the inverse, taking a known system and describing output behavior for given inputs. ("Simulate" is a specific method for making predictions, suggesting that there is a computational model of the system, complete at some level of detail.) CONTROL, not often associated with heuristic programs, takes a known system and determines inputs to generate prescribed outputs (Vemuri, 1978). Thus, these three operations, IDENTIFY, PREDICT, and CONTROL, logically cover the possibilities of problems in which one factor of the set {input, output, system} is unknown.

Both MONITOR and DIAGNOSE presuppose a pre-existing system design against which the behavior of an actual, "running" system is compared. Thus, one identifies the system with respect to its deviation from a standard. In the case of MONITOR, one *detects* discrepancies in behavior (or simply characterizes the current state of the system). In the case of DIAGNOSE, one *explains* monitored behavior in terms of discrepancies between the actual (inferred) design and the standard system.

To carry the analysis further, we compare our proposed terms to those used in Building Expert Systems:

- "Interpretation" is adopted as a generic category that broadly means to describe a working system. The most rudimentary form is simply identifying some unknown system from its behavior. Note that an identification strictly speaking involves a specification of constraints under which the system operates and a design (structure/process model). In practice, our understanding may not include a full design, let alone the constraints it must satisfy (consider the metarules of HERACLES (Section 6.1) versus our vague understanding of why they are reasonable). Examples of programs that identify systems are:
 - DENDRAL: system = molecular (structure) configuration (Buchanan, et al., 1969) (given spectrum behavior of the molecule).
 - PROSPECTOR: system = geological (formation) configuration (Hart, 1977) (given samples and geophysics behavior).
 - DEBUGGY: system = knowledge (program) configuration of student's subtraction facts and procedure (Burton, 1982) (given behavior on a set of subtraction problems).
- "Prediction" is adopted directly. Note that prediction, specifically simulation, may be an important technique underlying *all* of the other operations (e.g., using simulation to generate and test possible diagnoses).

- "Diagnosis" is adopted directly as a kind of IDENTIFICATION, with some part of the design characterized as faulty with respect to a preferred model.
- "Design" is taken to be the general operation that embraces both a characterization of structure (CONFIGURATION) and process (PLANNING).
- "Monitoring" is adopted directly as a kind of IDENTIFICATION, with system behavior checked against a preferred, expected model.
- "Debugging" is dropped, deemed to be equivalent to DIAGNOSIS plus MODIFY.
- "Repair" is more broadly termed MODIFY; it could be characterized as transforming a system to effect a redesign, usually prompted by a diagnostic description. MODIFY operations are those that change the *structure* of the system, for example, editing a program or using drugs (or surgery) to change a living organism. Thus, MODIFY is a form of "reassembly" given a required design modification.
- The idea of "executing a plan" is moved to the more general term ASSEMBLE, meaning the physical construction of a system. DESIGN is conceptual; it describes a system in terms of spatial and temporal interactions of components. ASSEMBLY is the problem of actually putting the system together in the real world. For example, contrast R1's problem with the problem of having a robot assemble the configuration that R1 designed. ASSEMBLY is equivalent to planning at a different level, that of a system that builds a designed subsystem.
- "Instruction" is dropped because it is a composite operation that doesn't apply to every system. In a strict sense, it is equivalent to MODIFY.


In addition to the operations already mentioned, we add SPECIFY—referring to the separable operation of constraining a system description, generally in terms of interactions with other systems and actual realization in the world (resources affecting components). Of course, in practice design difficulties may require modifying the specification, just as assembly may constrain design (commonly called "design for manufacturing").

5.4. Configuration and planning

The distinction between configuration and planning requires some discussion. We will argue that they are two points of view for the single problem of designing a system. For example, consider how the the problem of devising a vacation travel plan is equivalent to configuring a

complex system consisting of travel, lodging, restaurant, and entertainment businesses and specifying how that system will service a particular person or group of people. "Configure" views the task as that of organizing objects into some system that as a functioning whole will process/transform another (internal) system (its input). "Plan," as used here, turns this around, viewing the problem in terms of how an entity is transformed by its interactions with another (surrounding) system. Figure 5-3 illustrates these two points of view.

Configuration



Planning




Figure 5-3: The design problem seen from two perspectives

VLSI design is a paradigmatic example of the "configuration" point of view. The problem is to piece together physical objects so that their behaviors interact to produce the desired system behavior.

The "planning" point of view itself can be seen from two perspectives depending on whether a subsystem or a surrounding global system is being serviced:

1. *We service some system by moving it around for processing by subsystems of a surrounding world.* Paradigmatic examples are experiment planning (e.g., MOLGEN (Stefik, 1980, Friedland, 1979)) and shop scheduling (e.g., ISIS (Fox and Smith, 1984)). ASSEMBLY always involves planning of this form, and strictly speaking Stefik's MOLGEN solves an assembly problem, designing a system that physically constructs a DNA/cell configuration (a pre-designed subsystem). Equivalent examples are errand planning (Hayes-Roth and Hayes-Roth, 1979), vacation, and education plans. Here there is a well-defined object that is transformed by a well-defined sequence of interactions. In general, we do not care how the surrounding system is modified by these interactions, except that there are resource constraints affecting planning when many systems are being serviced.

2. *We specify how a well-defined object system will service a larger system in which it is contained.* Servicing is done by "moving the object system around." The paradigmatic example is the traveling salesman problem. Most realistic problems are hybrid because the "service subsystem" is resource limited and must be "restocked" and "refueled" by the surrounding system. *Truckin'*, the game used for teaching LOOPS (Stefik, et al., 1983), makes this clear. The traditional traveling salesman problem takes this form when allowance is made for food or fuel stops, etc.

While we appear to have laid out three perspectives on design, they are all computationally equivalent. It's our point of view about *purpose* and structuredness of interactions that makes it easier to understand a system in one way rather than another. In particular, in the first form of planning, the serviced subsystem is getting more organized as a result of its interactions. The surrounding world is modified in generally entropy-increasing ways as its resources are depleted. In the second form of planning the serviced world is getting more organized, while the servicing subsystem depletes its resources. Without considering the entropy change, *there is just a single point of view of a surrounding system interacting with a contained subsystem.*

"Configuration" is concerned with the construction of well-structured systems. In particular, if subsystems correspond to physically-independent components, design is equivalent to organizing pieces so they spatially fit together, with flow from input to output ports producing the desired result. (Note that R1 is given some of the pieces, not the functional properties of the computer system it is configuring. The functional design is implicit in the roster of pieces it is asked to configure—the customer's order.) It is a property of any system that can be described in this way that it is hierarchically decomposable into modular, locally interacting subsystems—the definition of a well-structured system. As Simon (Simon, 1969) points out, it is sufficient for design tractability for systems to be "nearly decomposable," with weak, but non-negligible interactions among modules.

Now, to merge this with the conception of "planning," consider how an abstract process can be visualized graphically in terms of a diagram of connected operations. The recent widespread availability of computer graphics has revolutionized how we visualize systems (processes and computations). Examples of traditional and more recent attempts to visualize the structure of processes are:

- *Flowcharts.* A program is a system. It is defined in terms of a sequence of operations for transforming a subsystem, the data structures of the program. Subprocedures and sequences of statements are subsystems that are structurally blocked and connected.
- *Automata theory.* Transition diagrams are one way of describing finite state

machines. Petri nets and dataflow graphs are other, related, notations for describing computational processes (see (Sowa, 1984) for discussion).

- *Actors*. A system can be viewed in terms of interacting, independent agents that pass messages to one another. Emphasis is placed on rigorous, local specification of behaviors (Hewitt, 1979). Object-oriented programming (Goldberg and Robson, 1983) is in general an attempt to characterize systems in terms of a configuration, centering descriptions on objects that are pieced together, as opposed to centering on data transformations.
- *Thinglab*. (Borning, 1979) emphasized the use of multiple, graphic views for depicting a dynamic simulation of mutually constrained components of a system. Borning mentions the advantages of visual experimentation for understanding complex system interactions.
- *Rocky's Boots*. In this personal computer game⁹, icons are configured to define a program, such as a sorting routine that operates on a conveyor belt. Movement icons permit automata to move around and interact with each other, thus describing "planning" (how systems will interact) from a "configuration" (combination of primitive structures) point of view.
- *FLIPP Displays*. Decision rules can be displayed in analog form as connected boxes that are interpreted by top-down traversal (see Figure 5-4). Subproblems can be visually "chunked"; logical reasoning can be visualized in terms of adjacency, blocking, alternative routes, etc. Characteristic of analog representations, such displays are economical, facilitating encoding and perception of interactions (Mackinlay and Genesereth, 1984).
- *Streams*. The structure of procedures can be made clearer by describing them in terms of the *signals* that flow from one stage in a process to another (Abelson, et al., 1985). Instead of modeling procedures in terms of time-varying objects (variables, see "planning" in Figure 5-3), we can describe procedures in terms of time-invariant streams. For example, a program might be characterized as ENUMERATE + FILTER + MAP + ACCUMULATE, a configuration of connected subprocesses. Stream descriptions, inspired by signal processing diagrams, allow a programmer to visualize processes in a succinct way that reveals the structural

⁹The Learning Company, Menlo Park, CA.

similarity of programs.




Figure 5-4: Simple FLIPP display, encoding rules $A \rightarrow B$, $A \rightarrow C$, and $D \rightarrow C$. (From (Cox, 1984).)

These examples suggest that we have not routinely viewed "planning" problems in terms of system "configuration" because we have not had adequate notations for visualizing interactions. In particular, we have lacked tools for graphically displaying hierarchical interactions and movement of subsystems through a containing system. Certainly a large part of the problem is that interactions can be *opportunistic*, so the control strategy that affects servicing (in either form of planning) is not specifiable as a fixed sequence of interactions. The inability to graphically illustrate flexible strategies was one limitation of the original Actors formalism (Hewitt, 1979). On the other hand, control strategies themselves may be specifiable as a hierarchy of processes, even though they are complex and allow for opportunism. The representation of procedures in HERACLES (Section 6.1) as layered rule sets (corresponding to tasks) (with both data-directed reasoning encoded as a separate set of tasks and inherited "interrupt" conditions) is an example of a well-structured encoding of an opportunistic strategy. More generally, strategy might be graphically visualized as layers of object-level operations and agenda-processing operations.

In general, a configuration point of view is impossible when physical or planning structures are unstable, with many global interactions (Hewitt, 1979). It is difficult or impossible to plan in such a world; this suggests that most practical planning problems can be characterized in terms of configuration. It is interesting to note that replacing state descriptions (configurations) with process descriptions has played an important role in scientific understanding of the origins of systems (Simon, 1969). As illustrated by the examples of this section, to understand these processes, we return to a configuration description, but now at the level of the structure of the process instead of the system it constructs or interprets.

5.5. Combinations of system problems

Given the above categorization of construction and interpretation problems, it is striking that expert systems tend to solve a sequence of problems pertaining to a given system in the world. Two sequences that commonly occur are:

- The Construction Cycle: SPECIFY + DESIGN {+ ASSEMBLE}

An example is R1 with its order processing front-end, XSEL. Broadly speaking,

selecting a book for someone in GRUNDY is *single-step planning*; the person is "serviced" by the book. Other examples are selecting a wine or a class to attend. The common sequence of terms in business, "plan and schedule," are here named SPECIFY (objectives) and PLAN (activities).

- **The Maintenance Cycle: {MONITOR + PREDICT +} DIAGNOSE + MODIFY**

This is the familiar pattern of medical programs, such as MYCIN. The sequence of MONITOR and PREDICT is commonly called *test* (repeatedly observing system behavior on input selected to verify output predictions). MODIFY is also called *therapy*.

This brings us back to the BES table (Figure 5-1), which characterizes INSTRUCTION and CONTROL as a sequence of primitive system operations. We can characterize the expert systems we have studied as such sequences of operations:

- MYCIN = MONITOR (patient state) + DIAGNOSE (disease category) + IDENTIFY (bacteria) + MODIFY (body system or organism)
- GRUNDY = IDENTIFY (person type) + PLAN (reading plan)
- SACON = IDENTIFY (structure type) + PREDICT (approximate numeric model) + IDENTIFY (classes of analysis for refined prediction)
- SOPHIE = MONITOR (circuit state) + DIAGNOSE (faulty module/component)

When a problem solver uses heuristic classification for multiple steps, as in GRUNDY, we say that the problem-solving method is *sequential heuristic classification*. Solutions for a given classification (e.g., stereotypes of people) become data for the next classification step. Note that Mycin does not strictly do sequential classification because it does not have a well-developed classification of patient types, though this is a reasonable model of how human physicians reason. However, it seems fair to say that MYCIN does perform a monitoring operation in that it requests specific information about a patient to characterize his state; this is clearer in NEOMYCIN and CASNET where there are many explicit, intermediate patient state descriptions. On the other hand, SOPHIE provides a better example of monitoring because it interprets global behavior, attempting to detect faulty components by comparison with a standard, correct model (discrepancies are violated assumptions).

It should be noted that how a problem is characterized in system terms may depend on our purpose, what "problem" we are attempting to solve in doing the system analysis or in building an expert system. For example, the OCEAN program (a product of Teknowledge, Inc.) checks configurations of computer systems. From a broad perspective, it is performing a MONITOR operation of the system that includes the human designer. Thus, OCEAN's inputs are the

constraints that a designer should satisfy, plus the output of his designing process. However, unlike DEBUGGY, we are not interested in understanding and correcting the designer's reasoning. Our purpose is to determine whether the computer system design meets certain specification constraints (e.g., power and space limitations) and to make minor corrections to the design. Thus, it seems more straightforward to say that OCEAN is doing a CONFIGURATION task, and we have given it a possible solution to greatly constrain its search.

Finally, for completeness, we note that robotics research is concerned chiefly with ASSEMBLY. Robotics is also converting CONTROL of systems from a purely numeric to a symbolic processing task. PREDICT in systems analysis has also traditionally involved numeric models. However, progress in the area of qualitative reasoning (also called mental models) (Bobrow, 1984) has made this another application for heuristic programming. Speech understanding is a strange case of identifying a system interaction between two speakers, attempting to characterize its output given a partial description (at the level of sounds) and environmental input (contextual) information.

Heuristic classification is particularly well-suited for problems of interpretation involving a system that is known to the problem solver. In this case, the problem solver can select from a set of systems he knows about (IDENTIFY), known system states (MONITOR), known system faults (DIAGNOSE), or known system behaviors (PREDICT/CONTROL). The heuristic classification method relies on experiential knowledge of systems and their behaviors. In contrast, constructing a new system requires construction of new structures (new materials or new organizations of materials). Nevertheless, we intuitively believe that experienced problem solvers construct new systems by modifying known systems. This confluence of classification and constructive problem solving is another important area for research.

Another connection the reader may have noticed: We made progress in understanding what expert systems do by describing them in terms of inference-structure diagrams. This vividly demonstrates the point made about streams, that it is highly advantageous to describe systems in terms of their configuration, *structurally*, providing dimensions for comparison. Gentner points out (Gentner and Stevens, 1983), that structural descriptions lie at the heart of analogy formation. A structural map of systems reveals similar relations among components, even though the components and/or their attributes may differ. This idea has been so important in research in the humanities during this century that it has been characterized as a movement with a distinct methodology, termed *structuralism* (De George and De George, 1972). The quotation by Bruner at the front of this paper describing the advantage of classification for a problem solver, applies equally well to the knowledge engineer.

6. INFERENCE STRATEGIES FOR HEURISTIC CLASSIFICATION

The arrows in inference-structure diagrams indicate the flow of inference, from data to conclusions. However, the actual order in which assertions are made is often not strictly left to right, from data to conclusions. This process, most generally called *search* or *inference control* has several aspects in heuristic classification:

- How does the problem solver get data? Is it supplied or must it be requested?
- If data is requested, how does the problem solver order his requests? (Called a *question-asking strategy*.)
- Does the problem solver *focus* on alternative solutions, requesting data on this basis?
- When new data is received, how is it used to make inferences?
- If there are choices to be made, alternative inference paths, how does the problem solver select which to attempt or which to believe?

In this section we first survey some well-known issues of focusing including data gathering, hypothesis testing, and data-directed inference. In this context, we introduce the HERACLES program, which is designed to solve problems by heuristic classification, and discuss its inference strategies. After this, we consider a kind of heuristic classification, termed *causal-process classification*, in order to understand the problem of choosing among inference paths. This discussion finally serves as a bridge to a consideration of non-classification or what we call *constructive* problem solving.

6.1. Focusing in heuristic classification

Focusing concerns what inferences the problem solver makes given new information or what inferences he attempts to make towards finding a solution.

The idea of a "triggering" relation between data and solutions is pivotal in almost all descriptions of heuristic classification inference (see (Rubin, 1975), (Szolovits and Pauker, 1978), (Aikins, 1983)). It is called a *constrictor* by Pople in recognition of how it sharply narrows the set of possible solutions (Pople, 1982). We say that "a datum triggers a solution" if the problem solver immediately thinks about that solution upon finding out about the datum. However, the assertion may be conditional (leading to an immediate request for more data) and is always context-dependent (though the context is rarely specified in our restricted-domain programs) (Clancey, 1984a). A typical trigger relation (from NEOMYCIN) is "Headache and red painful eye suggests glaucoma"—red, painful eye will trigger consideration of headache

and thus glaucoma, but headache alone will not trigger this association. In PIP (Pauker et al., 1976), there is a two-stage process in which possible triggers are first brought into working memory by association with solutions already under consideration. In general, *specificity*—the fact that a datum is frequently associated with just a few solutions—determines if a datum triggers a solution concept ("brings it to mind") in the course of solving a problem.

Triggers allow search to non-exhaustively combine reasoning forwards from data and backwards from solutions. Simple classification is constrained to be hierarchically top-down or directly bottom up from known data, but heuristic triggers make search opportunistic. Briefly, a given heuristic classification network of data and solution hierarchies might be interpreted in three ways:

1. *Data-directed search*: The program works forwards from data to abstractions, matching solutions until all possible (or non-redundant) inferences have been made.
2. *Solution- or Hypothesis-directed search*: The program works backwards from solutions, collecting evidence to support them, working backwards through the heuristic relations to the data abstractions and required data to solve the problem. If solutions are hierarchically organized, then categories are considered before direct features of more specific solutions.
3. *Opportunistic search*: The program combines data and hypothesis-directed reasoning (Hayes-Roth and Hayes-Roth, 1979). Data abstraction rules tend to be applied immediately as data become available. Heuristic rules trigger hypotheses, followed by a focused, hypothesis-directed search. New data may cause refocusing. By reasoning about solution classes, search need not be exhaustive.

Data- and hypothesis-directed search are not to be confused with the implementation terms "forward" or "backward chaining." R1 provides a superb example of how different the implementation and knowledge-level descriptions can be. Its rules are *interpreted* by forward-chaining, but it does a form of hypothesis-directed search, systematically setting up subproblems by a fixed procedure that focuses reasoning on spatial subcomponents of a solution (McDermott, 1982).

The degree to which search is focused depends on the level of indexing in the implementation and how it is exploited. For example, MYCIN's "goals" are solution classes (e.g., types of bacterial meningitis), but selection of rules for specific solutions (e.g., E.coli meningitis) is unordered. Thus, MYCIN's search within each class is unfocused (Clancey, 1983b). Generalized heuristics, of the form "data class implies solution class" (e.g., "compromised host implies gram-negative rods" or "flowers imply underlying rocks") make it possible to focus search on useful heuristics in both directions (e.g., if looking for serpentine rock, recall that

flowers identify rocks; if describing area and see flowers, recall that flowers identify rocks).¹⁰

Opportunistic reasoning requires that at least some heuristics be indexed so that they can be applied in either direction, particularly so new data and hypothesized solutions can be related to previously considered data and solutions. The HERACLES program (Heuristic Classification Shell, the generalization of NEOMYCIN) cross-indexes data and solutions in several ways that were not available in EMYCIN. HERACLES' inference procedure consists of 75 metarules comprising 40 reasoning tasks (Clancey, 1984a). Focusing strategies include:

- working upwards in type hierarchies before gathering evidence for subtypes;
- discriminating hypotheses on the basis of their descriptions as processes;
- making inferences that relate a new hypothesis to previously received data;
- seeking general classes of data before subclasses; and
- testing hypotheses by first seeking triggering data and necessary causes.

In HERACLES, the operators for making deductions are abstract, each represented by a set of metarules, corresponding to a procedure or alternative methods for accomplishing a task. Such a representation makes the explicit the inference control that is implicit in MYCIN's rules (Clancey, 1983b). As an example, Figure 6-1 illustrates how NEOMYCIN's abstract operators use backward deduction to confirm a hypothesized solution. (The program attempts to test the hypothesis TB by applying a domain rule mentioning ocular nerve dysfunction; to find this out, the program attempts to rule it out categorically, discovering that there is a CNS problem, but there are no focalsigns; consequently, the domain rule fails.)

As a second example, the following are some of the forward-deductive data-interpretation operators that HERACLES uses to relate a new datum to known solutions.

- Finding out more specific information so that a datum can be usefully related to hypotheses (e.g., given that the patient has a headache, finding out the duration and location of the headache).

¹⁰How human knowledge is indexed plays a major role in knowledge acquisition dialogues. The heuristic-classification model suggests that it may be efficient to proceed from data classes, asking the expert for associated solution classes. But it may be difficult to enumerate data classes. Instead, the expert might find it easier to work backwards from solutions (e.g., book categories) and then use a generate and test method to scan data prototypes (e.g., people stereotypes) for a match. Knowledge acquisition for heuristic classification is briefly considered in (Clancey, 1984b). See also the discussion of ETS in Section 10.

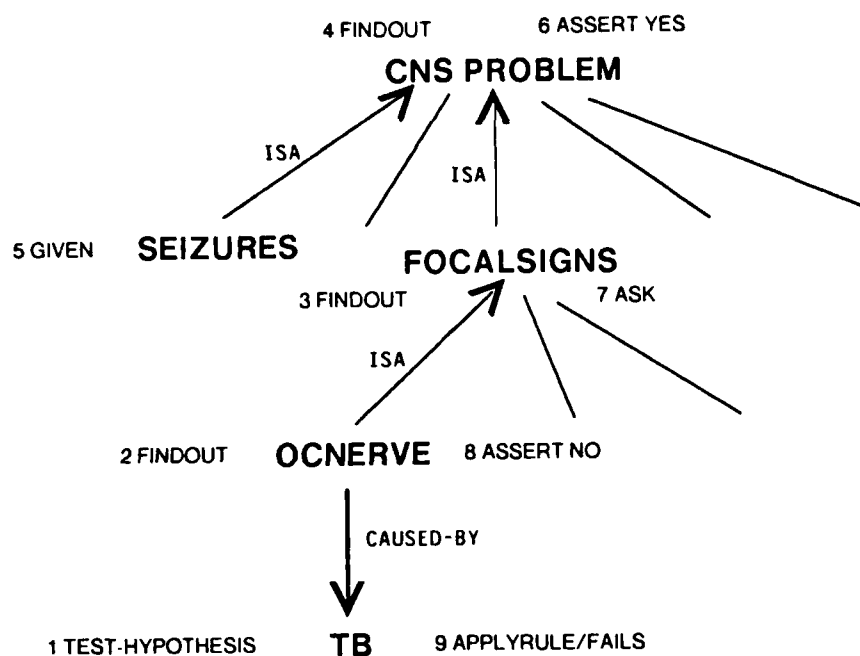


Figure 6-1: Backward deduction in NEOMYCIN to confirm a solution

- Making deductions that use the new datum to confirm "active" solutions (those previously considered, their taxonomic ancestors, and immediate siblings), sometimes called "focused forward-reasoning."
- Triggering possible solutions (restricted to abnormal findings that must be explained or "non-specific" findings not already explained by active solutions).

In general, the rationale for an inference procedure might be very complex. A study of HERACLES' inference procedure reveals four broad categories of constraints:

- mathematical (e.g., efficiency advantages of hierarchical search)
- sociological (e.g., cost for acquiring data),
- cognitive (e.g., computational resources), and
- problem expectations (e.g., characteristics of typical problems).

These are discussed in some detail in (Clancey, 1984a). Representing inference procedures so they can be explained and easily modified is currently an important research topic (e.g., see

(Clancey, 1983a, Genesereth, 1983, Neches, et al., 1985). Making the *assumptions* behind these procedures explicit so they can be reasoned about and dynamically modified is a challenging issue that AI is just beginning to consider.

In its inference procedure representation, HERACLES brings together the advantages of rule and frame representations. The "frame" paradigm advocates the point of view that domain concepts can be represented as hierarchies, separated from the inference procedure—an essential point on which the generality of the heuristic classification model depends. On the other hand, the "rule" paradigm demonstrates that much of the useful problem-solving knowledge is in non-hierarchical associations, and that there are clear engineering benefits for procedures to be encoded explicitly, as well-indexed conditional actions. In HERACLES domain concepts are hierarchically related; domain rules represent heuristic, non-hierarchical associations; and metarules represent an inference procedure that interprets the domain knowledge, solving problems by heuristic classification. The architecture of HERACLES, with details about the encoding of metarules in MRS (Genesereth et al., 1981), the metarule compiler, and explanation program are described in (Clancey, 1985).

6.2. Causal-process classification

A generic form of heuristic classification, commonly used for solving diagnostic problems, is *causal-process classification*. Data are generally observed malfunctions of some device, and solutions are abnormal processes causing the observed symptoms. We say that the inferred model of the device, the diagnosis, *explains* the symptoms. In general, there may be multiple causal explanations for a given set of symptoms, requiring an inference strategy that does not realize every possible association, but must reason about *alternative chains of inference*. In the worst case, even though diagnostic solutions are pre-enumerated (by definition), assertions might be taken back, so reasoning is non-monotonic. However, the most well-known programs that solve diagnostic problems by causal-process classification are monotonic, dealing with alternative lines of reasoning by assigning weights to the paths. Indeed, many programs do not even compare alternative explanations, but simply list all solutions, rank-ordered.

In this section, we will study SOPHIE in more detail (which reasons non-monotonically, using assumption-based belief maintenance), and compare it to CASNET (Weiss, et al., 1978) (which compares alternative chains of inference without explaining contradictions), and NEOMYCIN (Clancey and Letsinger, 1984) (which reasons exhaustively, using certainty factors to rank alternative inference chains). In these programs, solutions are pre-enumerated, but paths to them must be constructed. Our study serves several purposes: 1) to use the heuristic classification model to relate electronic and medical diagnosis, revealing that medical programs are generally trying to solve a broader problem; 2) To describe alternative heuristic classification inference strategies; and 3) To distinguish between classification and constructive problem solving.

6.2.1. Electronic and medical diagnosis compared

In SOPHIE, valid and abnormal device states are exhaustively enumerated, can be directly confirmed, and are causally related to component failures. None of this is generally possible in medical diagnosis, nor is diagnosis in terms of component failures alone sufficient for selecting therapy. Medical programs that deal with multiple disease processes (unlike MYCIN) do reason about abnormal states (called *pathophysiologic states*, e.g., "increased pressure in the brain"), directly analogous to the abnormal states in SOPHIE. But curing an illness generally involves determining the cause of the component failure. These "final causes" (called diseases, syndromes, etiologies) are processes that affect the normal functioning of the body (e.g., trauma, infection, toxic exposure, psychological disorder). Thus, medical diagnosis more closely resembles the task of computer system diagnosis in considering how the body relates to its environment (Lane, 1980). In short, there are two problems: First to explain symptoms in terms of abnormal internal states, and second to explain this behavior in terms of external influences (as well as congenital and degenerative component flaws). This is the inference structure of programs like CASNET and NEOMYCIN (Figure 6-2).

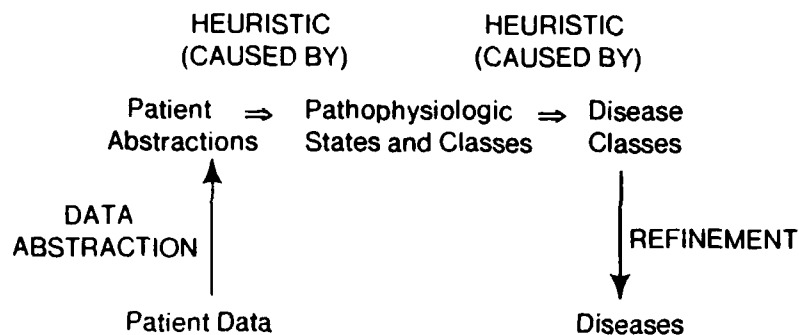


Figure 6-2: Inference structure of causal process classification

A network of pathophysiologic states causally relates data to diseases. States are linked to states, which are then linked to diseases in a classification hierarchy. Diseases may also be non-hierarchically linked by heuristics ("X is a complication of Y" (Szolovits and Pauker, 1978)). The causal relations in these programs are heuristic because they assume certain physiologic structure and behavior, which is often poorly understood and not represented. In contrast with pathophysiologic states, diseases are abstractions of *processes*—causal stories with agents, locations, and sequences of events. Disease networks are organized by these process features (e.g., an organ system taxonomy organizes diseases by location). A more general term for disease is *disorder stereotype*. In *process control* problems, such as chemical

manufacturing, the most general disorder stereotypes correspond to stages in a process (e.g., mixing, chemical reaction, filtering, packaging). Subtypes correspond to what can go wrong at each stage (Clancey, 1984a).

Structure/function models are often touted as being more general by doing away with "ad hoc symptom-fault rules" (Genesereth, 1984). But programs that make a single fault assumption, such as DART, select a diagnosis from a pre-enumerated list, in this case negations of device descriptions, e.g., (NOT (XORG X1)), "X1 is not an exclusive-or gate." However, a structure/function model makes it possible to *construct tests* (see Section 7). Note that it is not generally possible to construct structure/function models for the human body, and is currently even impractical for the circuit SOPHIE diagnoses (IP-28 power supply).

To summarize, a knowledge-level analysis reveals that medical and electronic diagnosis programs are not all trying to solve the same kind of problem. Examining the nature of solutions, we see that in an electronic circuit diagnosis program like SOPHIE solutions are component flaws. Medical diagnosis programs like CASNET attempt a second step, *causal process classification*, which is to explain abnormal states and flaws in terms of processes external to the device or developmental processes affecting its structure. It is this experiential knowledge—what can affect the device in the world—that is captured in disease stereotypes. This knowledge can't simply be replaced by a model of device structure and function, which is concerned with a different level of analysis.

The heuristic classification model and our specific study of causal process classification programs significantly clarifies what NEOMYCIN knows and does, and how it might be improved:

1. Diseases are processes affecting device structure and function;
2. Disease knowledge takes the form of schemas;
3. Device history schemas (classes of patients) are distinct from diseases;
4. Pathophysiologic states are malfunctioning module behaviors.

Furthermore, it is now clear that the original (bacteremia) MYCIN system does a combination of heuristic classification (using patient information to classify cultures as contaminated or significant) and simple classification (matching features of organisms, such as Gram stain and morphology). The meningitis knowledge base is more complex because it can infer the organism class heuristically, from patient abstractions, without having a culture result. NEOMYCIN goes a step further by dealing with *different processes* (infectious, trauma, psychogenic, etc.) and reasoning backwards from internal descriptions of current system states (e.g., brain mass lesion) to determine original causes (etiologies).

An important idea here is that medical diagnostic programs should separate descriptions of people from descriptions of diseases heuristically associated with them. Triggers should suggest patient types, just as they select diseases. Thus, medical diagnostic reasoning, when it takes the form of heuristic classification, is analogous to the problem-solving stages of GRUNDY, the expert librarian.

6.2.2. Inference control for coherency

As mentioned above, programs differ in whether they treat pathophysiologic states as independent solutions (NEOMYCIN) or find the causal path that best accounts for the data (CASNET). If problem features interact, so that one datum causes another ($D1 \rightarrow D2$ in Figure 6-3), then paths of inference cannot be correctly considered independently. The second feature explains the first, so classifications (alternative explanations) of the former can be omitted; there is a "deeper cause" ($C2$ dominates $C1$). This presumes a single fault, an assumption common to programs that solve problems by heuristic classification. CASNET uses a more comprehensive approach of finding the path with the greatest number of confirmed states and no denied states. The path describes a causal process, with an etiology or ultimate cause at the head and the path of states linking the etiology to the findings serving as a causal explanation of the findings.




Figure 6-3: Interacting data in classification

In the simplest rule-based systems, such as MYCIN, search is exhaustive and alternative reasoning paths are compared by numerical belief propagation (e.g., certainty factors). For example, Figure 6-4 shows that a datum, $D1$, is explained by two processes, $C1$ and $C2$. MYCIN and NEOMYCIN would make all three inferences, using certainty factor combination to determine which of $C1$ and $C2$ is more likely.

A more complicated approach involves detecting that one reasoning path is subsumed by another, such as the conflict-resolution strategy of ordering production rules according to specificity. HERACLES/NEOMYCIN's treatment of non-specific and "red flag" triggers is similar. In this case, assuming that $D1$ is a non-specific finding (associated with many disorders and may not need to be explained) and $D2$ is a red-flag finding (a serious symptom that must be explained) that triggered $C2$, NEOMYCIN will not make the inference relating $D1$ to $C1$ because $D1$ is already explained by $C2$. Therefore, $C1$ will not be added to the list of possible




Figure 6-4: Multiple explanations for a datum

solutions.

After finding some single classification that accounts for the most data, a residue of unexplained findings may remain. The approach used in CASNET and INTERNIST is to remove the explained data and simply begin another cycle of classification to explain what remains. In our example, both D1 and D2 would be explained by C2, so nothing would remain.

To summarize, when there are multiple causal links for classifying data—multiple explanations—inference must be controlled to avoid redundancy, namely multiple explanations where one would have been sufficient. The aim is to produce a *coherent* model that is *complete* (accounting for the most data) and *simple* (involving one fault process, if possible). In contrast with the idea of focussing discussed earlier, coherency places constraints on the sum total of inferences that have been made, not just the order in which they are made.

Of course, for an explanation based on a pathophysiological network to be coherent, it is necessary that inferences be consistent. For example, if D1 and D2 are contradictory, the network shown in Figure 6-4 will not produce consistent explanations. (C2 would depend on contradictory facts.) Presumably the knowledge engineer has ensured that all paths are consistent and that contradictory alternatives are explicit (e.g., by introducing (NOT D1) to the path including D2).

An ideal way to avoid these problems is to perform the diagnosis using a model of a correctly working device, in contrast with a network of pathophysiological states. This is the method used by SOPHIE. A consistent interpretation includes the observed data and assumptions about the operation of circuit components. A fault is detected by making inferences about circuit behavior in the case at hand until a contradiction is found. Specifically, SOPHIE uses assumption-based belief maintenance to detect faults. It propagates constraints (describing device behavior), records assumptions (about correct behavior of components and modules) upon which inferences are based, explains contradictory inferences in terms of violated assumptions, and makes measurements to narrow down the set of possibly violated assumptions to a single fault. Making assumptions explicit and reasoning about them ensures coherency, rather than relying on its implicit and ad hoc encoding in the design of a state network.

6.2.3. Multiple solutions and levels of detail

The first step beyond selecting single pre-enumerated faults is to dynamically construct *sets of alternative faults*, as proposed for CADUCEUS (Pople, 1982). Each set of faults constitutes a *differential diagnosis* or set of alternative diagnoses. Each diagnosis consists of one or more faults. A diagnosis of multiple faults is constructed by operators that combine disorders in terms of subtype and cause. For example, referring to Figure 6-4, one differential diagnosis would include C1 & C2; another would include C2, but not C1.

The next more complicated approach allows for interactions among disorders, as in ABEL (Patil, 1981a). Interaction can take the form of masking or subtracting (quantitative) effects, summation of effects, or superimposition with neither subtraction nor summation. These interactions are predicted and explained in ABEL by finding a state network, including normal states, that is *consistent on multiple levels of detail*. Combinatorial problems, as well as elegance, argue against pre-enumerating such networks, so solutions must be constructed. Each diagnostic hypothesis is a separately constructed *case-specific model*—the links describing the individual case do not all pre-exist in the knowledge base.

A simple way of comparing this kind of reasoning to what occurs in classification is to consider how concepts are instantiated and related in the two approaches. In a program like MYCIN, there are case-specific instances, but they are only the most general concepts in the knowledge base—the patient, laboratory cultures, organisms, and drugs. Links among these instances (or individuals), constituting the "context tree," are dynamically created (albeit given as data) to form a case-specific model. In contrast, in ABEL case-specific, constructed descriptions are also at the level of individual disorders and their causes—disease components are instantiated and linked together in novel ways (thus allowing for interaction among diseases).

6.2.4. Constructing implication paths vs. constructing solutions

If all programs construct inference paths, aren't they all solving problems by construction of solutions? At issue is a point of view about what is a solution.

In NEOMYCIN, CASNET, and SOPHIE, the solutions are single faults, pre-enumerated. Reasoning about inference paths is a mechanism for selecting these solutions. While an inference path through the causal network of CASNET or NEOMYCIN is a disease process description, it is only a linear path, no different from a chain of implication in MYCIN. While links represent subtype and cause, they are interpreted in a uniform way for propagating weights. We conclude that if there is only one operator for building inference paths, the program is not constructing a solution, but is only selecting the nodes at the end points of its reasoning chains. All of the programs we characterize as doing constructive problem solving either have a generator for solutions or must choose among multiple operators for constructing a solution. The solutions aren't explicitly enumerated, so there can be *no pre-existing links for mapping*

problem descriptions to solutions directly. In ABEL and CADUCEUS, solutions are descriptions of disease processes, constructed by operators for incrementally elaborating and aggregating solution components, which is more than just propagating belief (what we commonly call "implication"). The constructed solution is not simply an inference path from data to solution, but a *configuration* of primitive solution components; these programs *configure disease descriptions*, they do not select them directly.

It should now be abundantly clear that it is incorrect to say that diagnosis is a "classification problem." As Pople concluded in analyzing the inadequacies of INTERNIST, only *routine* medical diagnosis problems can be solved by classification (Pople, 1982). When there are multiple diseases, there are too many possible combinations for the problem solver to have considered them all before. The problem solver must construct a consistent network of interacting diseases to explain the symptoms. The problem solver *formulates a solution*; he doesn't just make yes-no decisions from a set of fixed alternatives. For this reason, Pople calls non-routine medical diagnosis an ill-structured problem (Simon, 1973), though it may be more appropriate to reserve this term for the theory formation task of the physician-scientist who is defining new diseases. To make the point most boldly: For GRUNDY, the librarian, to satisfy a reader by constructing a solution, she would have to write a new book.

7. CONSTRUCTIVE PROBLEM SOLVING, AN INTRODUCTION

In a study of problem solving, Greeno and Simon characterize kinds of problems in terms of the constraints imposed on the problem solver.

In a transformation problem such as the Tower of Hanoi, or finding a proof for a theorem, the goal is a specific [given] arrangement of the problem objects, such as a specific location for all of the disks in the Tower of Hanoi, or a specific expression to be proved in logic. Thus, the question is not what to construct, as it is in a problem of design, but how the goal can be constructed with the limited set of operators available.... (Greeno and Simon, 1984)

While different tasks do impose different constraints on the problem solver, we have argued that experiential knowledge allows a "design" problem to be solved as if it were a "transformation" problem. For while design problems may not generally provide the problem solver with a specific solution to attain, he may from experience know of a solution that will work. In heuristic classification the solution space is known to the problem solver as a *set* of explicit alternatives, and problem solving takes the form of "proving" that one of them is best or acceptable.¹¹

¹¹In adopting the heuristic classification model as a psychological theory, we must be more careful about this issue of explicitness. Human memory has properties different from knowledge base representations, so there is a difference between "explicitly known now" and "previously known." In practice, remembering a previous solution may require reconstruction, and hence some elements of constructive problem-solving.

For example, in diagnostic programs that assume a single fault, such as NEOMYCIN, CASNET, SOPHIE, and DART, the inference process is equivalent to finding the most specific and most likely theorem (solution) that can be proved correct. Thus, the spectrum of problem solving effort and methodology is aligned at least as much with experience as with the nature of the task. Amarel makes this point in distinguishing between "derivation" and "formation" problems (Amarel, 1978), emphasizing that experience provides knowledge for mapping problem conditions to solutions. Thus, experience moves tasks for a given problem solver to the "derivation" end of the spectrum—heuristic classification.

Often problems of DESIGN and DIAGNOSIS are not amenable to solution by heuristic classification because possible final "states" cannot be practically enumerated, exhaustively learned (from experience or direct teaching), or for some reason a previously used solution is just not acceptable; solutions must be constructed rather than selected. However, even when solutions are constructed, classification might play a role, for example, in planning the problem-solving process or in selecting the pieces to be configured.

The essential differences between heuristic construction and classification are the need for some "data structure" to post the assembled solution and operators for proposing and reasoning about solution fragments (Erman, et al., 1981). In classification, triggers focus the search, but may not be necessary; controlled forward-deduction from given data and hierarchical search in the set of fixed solutions may be sufficient. In construction, triggers may be essential, as well as knowledge about how parts of a solution fit together.

The following are some examples of heuristic construction operators:

- In HASP, the ocean surveillance signal interpretation system (Nii, et al., 1982), one operator attaches a new line segment (from the input sensor) to a previous line that was last heard less than thirty minutes ago (with a certainty of .5), thus extending the model of the location of a particular vessel.
- ABEL has six "structure building" operators, including *projection* (to hypothesize associated findings and diseases suggested by states in the case-specific model) and *causal elaboration* (to determine causal relations between states at a detailed level, based on causal relations between states at the next aggregate level) (Patil, 1981b).
- AM has operators for proposing (syntactic) structural modifications to concepts. For example, a concept is generalized by deleting a conjunction in its characteristic function definition in Lisp (Lenat and Brown, 1984). One lesson of Eurisko is that complex concept formation requires a more extensive set of operators (defined in terms of conceptual relations, or as Lenat puts it, "slots for describing heuristics").
- MOLGEN (Stefik, 1980) has both assembly-design operators and laboratory-domain

operators. PROPOSE-OPERATOR is a design operator that proposes a laboratory operator that will reduce differences between the goal and current state, extending the plan forward in time. It is "responsible for linking new laboratory steps correctly to the neighboring laboratory steps and goals." There are four kinds of physical, structure-modifying laboratory operators: merge, amplify, react, and sort.

- In DART, diagnosis is done by classification (and the use of the proof method is made explicit in the program description), but testing the circuit to gather more information (MONITOR) is done by construction. The abstract operator (IF (AND $a_1 \dots a_m$ Ob) THEN (OR (NOT P_1) ... (NOT P_n))) serves as a template for generating tests, where the a_i 's are achievable settings or structure changes, Ob is one or more observations to be made, and the P_i 's are assumptions about correct device structure or function. Thus, a particular device setting and observations will confirm that one or more assumptions are violated, narrowing down the set of possible faults (similar to SOPHIE). Note that the heuristics used in DART are general search strategies used to control the deductive process, not domain-specific links. DART has heuristics and uses classification (for diagnosis), but it does not do heuristic classification, in the form we have described it. Specifically, it lacks experiential, schema knowledge for classifying device states and describing typical disorders.
- MYCIN's antibiotic therapy program (Clancey, 1984c) generates combinations of drugs from "instructions" that abstractly describe how the number of drugs and their preference are related. These generator instructions can be viewed as operators (or a grammar) for constructing syntactically correct solutions.

To summarize the alternative means of computing solutions we have seen:

1. Solutions might be *selected* from a pre-enumerated set, by classification.
2. Solutions may be *generated* whole, as in DENDRAL and MYCIN's therapy program.
3. Solutions may be *assembled* from primitives, incrementally, as in HASP.

As Simon indicates (Simon, 1969), these methods can be combined, sequentially or hierarchically (as in hierarchical planning), with perhaps alternative decompositions for a single system.

8. RELATING TOOLS, METHODS, AND TASKS

In our discussion we have emphasized the question, "What is the method for computing a solution?" We have made a distinction between data and solution in order to clarify the large-scale computational issues of constructing a solution versus selecting it from a known set. The logical next step is to relate what we have learned about conceptual structures, systems problems, and the classification/construction distinction to the tools available for building expert systems. Conventionally, the thing to do at this point would be to provide a big table relating tools and features. These can be found in many books (e.g., (Harmon and King, 1985)). The new analysis we can provide here is to ask which tool features are useful for heuristic classification and which are useful for constructive problem solving.

While simple rule-based languages like EMYCIN omit knowledge-level distinctions we have found to be important, particularly schema hierarchies, they have nevertheless provided an extremely successful programming framework for solving problems by classification. Working backwards (backchaining) from a pre-enumerated set of solutions guarantees that only the relevant rules are tried and useful data considered. Moreover, the program designer is encouraged to use means-ends analysis, a clear framework for organizing rule writing. EMYCIN and other simple rule-based systems are appropriate for solving problems by heuristic classification because inference chains are short (commonly five or fewer steps between raw data and solution), and each new rule can be easily viewed as adding a link in the mapping between data (or some intermediate abstraction) and solutions.

With the advent of more complex knowledge representations, such as KL-ONE, it is unclear whether advantages for explicit representation will be outweighed by the difficulty of learning to use these languages correctly. The analysis needed for identification of classes and relations and proper adherence to representational conventions might require considerable experience, or even unusual analytical abilities (recall the analysis of concepts in Section 4.5). Recent research indicates that it might be difficult or practically impossible to design a language for conceptual structures that can be unambiguously and consistently used by knowledge engineers (Brachman, et al., 1983). Just as the rule notation was "abused" in MYCIN by ordering rule clauses to encode hierarchical and procedural knowledge, users of KL-ONE implicitly encode knowledge in structural properties of concept hierarchies, relying on the effect of the interpreter to make correct inferences. Brachman, et al. propose a model of *knowledge representation competence*, in which a program is told what is true and what it should do, and left to encode the knowledge according to its own conventions to bring about the correct

reasoning performance.¹²

While many of the conceptual structures and inference mechanisms required to encode a heuristic classification problem solver have now been identified, no knowledge engineering tool today combines these capabilities in a complete package. Perhaps the best system for classification we could imagine might be a combination of KL-ONE (so that conceptual relations are explicit and to provide automatic categorization of concepts (Schmolze and Lipkis, 1983)), HERACLES (so that the inference procedure is explicit, well-structured, and independent of domain knowledge representation), and SHRINK (Kolodner, 1982) (to provide automatic refinement of classifications through problem-solving experience). In this respect, it should be noted there is some confusion about the nature of heuristic classification in some recent commercial tools on the market. Close inspection reveals that they are capable of only simple classification, lacking structures for data abstraction, as well as a means to separate definitional features from heuristic associations between concepts (Harmon and King, 1985).

Regarding constructive problem solving, the major distinction among tools appears to be the method for coping with alternative choices in configuring a solution. Tools for constructive problem solving necessarily include methods for controlling search that go beyond the focusing operations found in tools that solve problems by classification. For example, well-known search control methods used in construction include: Least-commitment (Stefik, 1980) (avoiding decisions that constrain future choices), representing explicitly multiple "hypothetical" worlds (branching on choices to construct alternative solutions), variable propagation or relaxation (systematic refinement of solutions), backtracking (retracting constructions), version space search (Mitchell, 1982) (bounding a solution using variables and constraints), and debugging (Sussman, 1975) (modifying an unsatisfactory solution).

What have we learned that enables us to match problems to tools? Given a task, such as troubleshooting, we might have previously asked, "Is this tool good for diagnosis?" Now, we insert an intermediate question about *computational requirements*: Is it possible or acceptable to pre-enumerate solutions? Is it possible or acceptable to rank order solutions? Rather than matching tasks to tools directly, we interpose some questions about the method for computing

¹²See (Levesque, 1984) for details. In apparent conflict with our use of inference diagrams to describe what a heuristic-classification problem solver knows, Levesque says, "There is nothing to say about the *structure* of these abstract bodies of knowledge called knowledge bases." One way of resolving this is to say that knowledge *content* has structure, but knowledge-level specification is not *about* structures in the agent (problem solver). This is supported by Newell's remark, "Relationships exist between goals, of course, but these are not realized in the structure of the system, but in knowledge." (Newell, 1982) This is our intent in separating the abstract characterization of what a problem solver knows (heuristic classification model) from its encoding in the agent's symbol system (expert system representation).

solutions. The basic choice of classification versus construction is the missing link for relating implementation terminology ("rules," "blackboard," "units") to high-level conceptual structures and inference requirements.

In summary, we suggest the following sequence of questions for matching problems to tools:

1. Describe the problem in terms of a sequence of operations relating to systems. If the problem concerns ASSEMBLY or construction of a perceptual system, seek a specialist in another area of AI. If the problem concerns numerical PREDICTION or CONTROL, it might be solved by traditional systems analysis techniques.
2. Do constraints on customization or naturally occurring variety allow the solution space to be practically pre-enumerated? If so, use heuristic classification. If not, is there a hierarchical or grammatical description that can be used to generate possible solutions? Are there well-defined solution-construction operators that are constrained enough to allow an incremental (state-space) search?
3. Are there many uncertain choices that need to be made? If a few, exhaustive generation with a simple certainty-weighting model may be sufficient; if many, some form of lookahead or assumption/justification-based inference mechanism will be necessary.

The notation of inference-structure diagrams used in this paper can also be used to form a knowledge specification that can then be mapped onto the constructs of a particular knowledge representation language. First, identify and list possible solutions, data, and intermediate (more general) categories. Examining inference chains, classify links between concepts as definitional, type categorization, and heuristic. Then, draw an inference structure diagram to arrange relations within a type hierarchy vertically and show heuristics as horizontal lines. Finally, map this diagram into a given representation language. For example, subtype links are represented as ordered clauses in an EMYCIN rule: "If the patient has an infection, and the kind of infection is meningitis, and"

Our study suggests two additional perspectives for critiquing constructive tools. Viewing *solutions* as models of systems in the world, we require means for detecting and controlling the coherency (completeness and consistency) of inferences. In describing *computational methods* in terms of operators, we need means to construct, record, and relate inference graphs. We conclude that the method by which inference is controlled—how an inference graph representing a system model is computed—is a crucial distinction for comparing alternative knowledge engineering tools for constructive problem solving. Relating the above methods for constructing solutions (e.g., version space, least commitment, blackboard architecture) to problem tasks is beyond the scope of this paper. It is possible that the problem categories of Section 5 will be useful. Though they may prove to be an orthogonal consideration, as we

discovered in distinguishing between classification and construction.

9. KNOWLEDGE-LEVEL ANALYSIS

As a set of terms and relations for describing knowledge (e.g., data, solutions, kinds of abstraction, refinement operators, the meaning of "heuristic"), the heuristic classification model provides a *knowledge-level analysis* of programs (Newell, 1982). As defined by Newell, a knowledge-level analysis "serves as a specification of what a reasoning system should be able to do." Like a specification of a conventional program, this description is distinct from the representational technology used to implement the reasoning system. Newell cites Schank's conceptual dependency structure as an example of a knowledge level analysis. It indicates "what knowledge is required to solve a problem... how to encode knowledge of the world in a representation." It should be noted that Newell intends for the knowledge-level specification to include the closure of what the reasoning system *might* know. Our approach to this problem is to characterize the problem solver's computational method and the structure of his knowledge. What a heuristic classification problem solver "is able to do" is specified in terms of the patterns of problem situations and solutions he knows and the space of (coherent) mappings from raw data to solutions.

After a decade of "explicitly" representing knowledge in AI languages, it is ironic that the pattern of heuristic classification should have been so difficult to see. In retrospect, certain views were emphasized at the expense of others:

- *Procedureless languages.* In an attempt to distinguish heuristic programming from traditional programming, procedural constructs are left out of representation languages (such as EMYCIN, OPS, KRL (Lehnert and Wilks, 1979)). Thus, inference relations cannot be stated separately from how they are to be used (Hayes, 1977, Hayes, 1979).
- *Heuristic nature of problem solving.* Heuristic association has been emphasized at the expense of the relations used in data abstraction and refinement. In fact, some expert systems do only simple classification; they have no heuristics or "rules of thumb," the key idea that is supposed to distinguish this class of computer programs.
- *Implementation terminology.* In emphasizing new implementation technology, terms such as "modular" and "goal directed" were more important to highlight than the content of the programs. In fact, "goal directed" characterizes any rational system and says very little about how knowledge is used to solve a problem. "Modularity" is a representational issue of indexing, how the knowledge objects can be independently accessed.

Nilsson has proposed that logic should be the *lingua franca* for knowledge-level analysis (Nilsson, 1981). Our experience suggests that the value of using logic is in adopting a set of terms and relations for describing knowledge (e.g., kinds of abstraction). Logic is especially valuable as a tool for knowledge-level analysis because it emphasizes *relations*, not just implication.

10. RELATED ANALYSES IN PSYCHOLOGY AND ARTIFICIAL INTELLIGENCE

Only a monograph-length review could do justice to the vast amount of research that relates to heuristic classification. Every discipline from ancient philosophy through modern psychology seems to have considered some part of the story.

Several AI researchers have described heuristic classification in part, influencing this analysis. For example, in *CRYSTALIS* (Engelmore and Terry, 1979) data and solution abstraction are clearly separated. The *EXPERT* rule language (Weiss, 1979) similarly distinguishes between "findings" and a taxonomy of hypotheses. In *PROSPECTOR* (Hart, 1977), rules are expressed in terms of relations in a semantic network. In *CENTAUR* (Aikins, 1983), a variant of *EMYCIN*, solutions are explicitly *prototypes* of diseases. Chandrasekaran and his associates have been strong proponents of the classification model: "The normal problem-solving activity of the physician... (is) a process of classifying the case as an element of a disease taxonomy" (Chandrasekaran and Mittal, 1983, Gomez and Chandrasekaran, 1984). Recently, Chandrasekaran, Weiss, and Kulikowski have generalized the classification schemes used by their programs [*MDX* (Chandrasekaran, 1984) and *EXPERT* (Weiss and Kulikowski, 1984)] to characterize problems solved by other expert systems.

In general, rule-based research in AI emphasizes the importance of heuristic association; frame systems emphasize the centrality of concepts, schema, and hierarchical inference. A series of knowledge representation languages beginning with *KRL* have identified structured abstraction and matching as a central part of problem solving (Bobrow and Winograd, 1979). These ideas are well-developed in *KL-ONE*, whose structures are explicitly designed for classification (Schmolze and Lipkis, 1983).

Building on the idea that "frames" are not just a computational construct, but a theory about a kind of knowledge (Hayes, 1979), cognitive science studies have described problem solving in terms of classification. For example, routine physics problem solving is described by Chi (Chi, et al., 1981) as a process of data abstraction and heuristic mapping onto solution schemas ("experts cite the abstracted features as the relevant cues (of physics principles)"). The inference structure of *SACON*, heuristically relating structural abstractions to numeric models, is the same. In *NEWTON*, De Kleer referred to packages of equations, associated with problem features, as *RALCMs* (Restricted Access Local Consequent Methods) ("with this representation, only a few decisions are required to determine which equations are relevant") (de Kleer, 1979).

Related to the physics problem solving analysis is a very large body of research on the nature of schemas and their role in understanding (Schank, 1975, Rumelhart and Norman, 1983). More generally, the study of classification, particularly of objects, also called *categorization*, has been a basic topic in psychology for several decades (e.g., see the chapter on "conceptual thinking" in (Johnson-Laird and Wason, 1977) and (Rosch and Lloyd, 1978)). However, in

psychology the emphasis has been on the nature of categories and how they are formed (an issue of learning). The programs we have considered make an identification or selection from a pre-existing classification (an issue of memory retrieval). In recent work, Kolodner combines the retrieval and learning process in an expert system that learns from experience (Kolodner, 1982). Her program uses the MOPS representation, a classification model of memory that interleaves generalizations with specific facts (Kolodner, 1983).

Probably the most significant work on classification was done by Bruner and his colleagues in the 1950's (Bruner, et al., 1956). Bruner was concerned with the nature of concepts (or categories), how they were attained (learned), and how they were useful for problem solving. A few quotes illustrate the clarity and relevance of his work:

To categorize is to render discriminably different things equivalent, to group objects and events and people around us into classes, and to respond to them in terms of their class membership rather than their uniqueness. (Bruner, et al., 1956) (page 1)

...the task of isolating and using a concept is deeply imbedded in the fabric of cognitive life; that indeed it represents one of the most basic forms of inferential activity in all cognitive life. (page 79)

...[what] we have called "concept attainment" in contrast to "concept formation" is the search for and testing of attributes that can be used to distinguish exemplars from nonexemplars of various categories, the search for good and valid anticipatory cues. (page 233)

Bruner described some of the heuristic aspects of classification:

... regard a concept as a network of sign-significate inferences by which one goes beyond a set of observed criterial properties exhibited by an object or event in question, and then to additional inferences about other *unobserved* properties of the object or event. (page 244)

What does categorizing accomplish for the organism? ... it makes possible the sorting of functionally significant groupings in the world. (page 245)

We map and give meaning to our world by relating classes of events rather than by relating individual events. The moment an object is placed in a category, we have opened up a whole vista for "going beyond" the category by virtue of the superordinate and causal relationships linking this category to others. (page 13)

Bruner was well ahead of AI in realizing the centrality of categorization in problem solving. Particularly striking is his emphasis on strategies for selecting cues and examples, by which the problem solver directs his learning of new categories ("information gathering strategies"). Bruner's study of hypothesis formation and strategies for avoiding errors in learning is particularly well-developed, "For concern about error, we contend, is a necessary condition for evoking problem-solving behavior" (page 210) [compare to "failure-driven memory" (Schank, 1981) and "impasses" (VanLehn, 1983)].

On the other hand, Bruner's description of a concept is impoverished from today's point of view. The use of toy problems (colored cards or blocks, of course) suggested that categorization was based on "direct significant"—a logical combination of observable,

discriminating (perhaps probabilistic) features. This Aristotelian view persisted in psychology and psychometrics well into the 1970's, until the work of Rosch, who argues that concepts are prototypical, not sets of definitional features (Rosch, 1978, Mervis and Rosch, 1981, Cohen and Murphy, 1984, Greeno and Simon, 1984). Rosch's work was influenced by AI research, but it also had its own effect, particularly in the design of KRL (Bobrow and Winograd, 1979).

The heuristic classification model presented in this paper builds on the idea that categorization is not based on purely essential features, but rather is primarily based on *heuristic*, non-hierarchical, *but direct* associations between concepts. Bruner, influenced by game theory, characterizes problem solving (once a categorization was achieved) in terms of a payoff matrix; *cues are categorized to make single decisions of actions to take*. Influenced by the psychology of the day, he views problem solving in the framework of a stimulus and a response, where the stimulus is the "significant" (cues) and the response is the action that the categorization dictates. He gives examples of medical diagnosis in this form.

We have had the advantage of having a number of working models of reasoning to study—expert systems—whose complexity go well-beyond what Bruner was able to formally describe. We have observed that problem solving typically involves a sequence of categorizations. Each categorization can be characterized generically as an operation upon a system—specify, design, monitor, etc. Most importantly, we have seen that each classification is not a final consequence or objective in "a payoff matrix governing a situation" (page 239). Rather it is often a plateau chaining to another categorization. Bruner's payoff matrix encodes heuristic associations.

Building upon Rosch's analysis and developments in knowledge representation, recent research in cognitive science has significantly clarified the nature of concepts (Cohen and Murphy, 1984, Rosch, 1978). In particular, attention has turned to why concepts take the form they do. While many concepts are based on *natural kinds* (e.g., MYCIN's organisms and GRUNDY's books), others are *experiential* (e.g., reader and patient stereotypes of people), or *analytic* (e.g., SOPHIE's module behavior lattice and SACON's programs). Miller (Miller, 1978) suggests that formation of a category is partly constrained by its heuristic implication. Thus, therapeutic implication in medicine might serve to define diagnostic and person categories, working backwards from pragmatic actions to observables. This functional, even behavioral, view of knowledge is somewhat disturbing to those schooled in the definition of concepts in terms of essential features, but it is consistent with our analysis of expert systems. Future studies of what people know, and the nature of meaning, will no doubt depart even more from essential features to consider heuristic "incidental" associations in more detail.

Finally, learning of classifications has been a topic in AI for some time. Indeed, interest goes back to early work in pattern recognition. As Chandrasekaran points out (Chandrasekaran, 1984), it is interesting to conceive of Samuels' hierarchical evaluation

functions for checker playing as an implicit conceptual hierarchy. Recent work in classification, perhaps best typified by Michalski's research (Michalski and Stepp, 1983), continues to focus on learning essential or definitional features of a concept hierarchy, rather than heuristic associations between concepts. However, this form of learning, emphasizing the role of object attributes in classification, is an advance over earlier approaches that used numeric measures of similarity and lacked a conceptual interpretation. Also, working from the traditional research of psychometrics, Boose's ETS knowledge acquisition program (Boose, 1984) makes good use of a psychological theory of concept associations, called *personal construct theory*. However, ETS elicits only simple classifications from the expert, does not exploit distinctions between hierarchical, definition, and heuristic relations, and has no provision for data abstraction.

Perhaps the greatest value of the heuristic classification model is that it provides an overarching point of view for relating pattern recognition, machine learning, psychometrics, and knowledge representation research.

11. SUMMARY OF KEY OBSERVATIONS

The heuristic classification model may seem obvious or trivial after it is presented, but the actual confusion about knowledge engineering tools, problem-solving methods, and kinds of problems has been quite real in AI for the past decade. Some might say, "What else could it be? It had to be classification"—as if a magic trick has been revealed. But the point of this paper is not to show a new kind of trick, or a new way of doing magic tricks, but to demystify traditional practice.

Sowa's reference to Levi-Strauss' anthropological "systems analysis" is apt:

The sets of features ... seem almost obvious once they are presented, but finding the right features and categories may take months or years of analysis. The proper set of categories provides a structural framework that helps to organize the detailed facts and general principles of a system. Without them, nothing seems to fit, and the resulting system is far too complex and unwieldy.

Expert systems are in fact *systems*. To understand them better, we have given high-level descriptions of how solutions are computed. We have also related the tasks of these programs to the kinds of things one can do to or with a concrete system in the world. Below is a summary of the main arguments:

- A broad view of how a solution is *computed* suggests that there are two basic problem-solving methods used by expert systems: heuristic classification and construction.
- Kinds of inference in different stages of routine problem solving vary systematically, so data are often generalized or redefined, while solutions are more

often matched schematically and refined. A domain-specific heuristic is a direct, non-hierarchical association between different classes. It is not a categorization of a "stimulus" or "cue" that directly matches a concept's definition. Rather, there may be a chain of abstraction inferences before reaching categories that usefully characterize problem features. This pattern is shown two ways in Figures 2-3 and 4-5. The inference structure of heuristic classification is common in expert systems, independent of the implementation representation in rules, frames, ordinary code, or some combination.

- Selecting solutions involves a form of proof that is often characterized as derivation or planning; constructing solutions involves piecing together solutions in a manner characterized as configuration. To *select* a solution, the problem solver needs experiential ("expert") knowledge in the form of *patterns of problems and solutions* and heuristics relating them. To *construct* a solution, the problem solver applies models of structure and behavior, in the form of constraints and inference operators, by which objects can be designed, assembled, diagnosed, employed in some plan, etc.
- A broad view of kinds of problems, described in terms of synthesis and analysis of *systems*, suggests two points of view for describing a system's design: a configuration in terms of structural relations of functional components, versus a plan for the processes that characterize the system's behavior. From the point of view of a *system*, reasoning may involve a limited set of generic operations, e.g., MONITOR, DIAGNOSE, MODIFY. In heuristic classification, this takes the form of a sequence of mapping between classifications corresponding to each generic operation.
- In a manner analogous to *stream* descriptions of computer programs, the inference-structure diagrams used in this paper reveal the patterns of reasoning in expert systems.

12. IMPLICATIONS

A wide variety of problems can be solved by heuristic mapping of data abstractions onto a fixed, hierarchical network of solutions. This problem-solving model is supported by psychological studies of human memory and categorization. There are significant implications for expert systems research. The model provides:

- A *high-level structure for decomposing problems*, making it easier to recognize and represent similar problems. For example, problems can be characterized in terms of sequences of system classifications. Catalog selection (single-step planning)

programs might be improved by incorporating a more distinct phase of user modelling, in which needs or requirements are explicitly classified. Diagnosis programs might profitably make a stronger separation between device-history stereotypes and disorder knowledge. "Blackboard" systems might re-represent "knowledge sources" to distinguish between classification and construction inference operators.

- *A specification for a generic knowledge engineering tool* designed specifically for heuristic classification. The advantages for knowledge acquisition carry over into explanation and teaching.
- *A basis for choosing application problems.* For example, problems can be selected using the systems taxonomy (Figures 5-1 and 5-2), allowing knowledge engineers to systematically gain experience in different kinds of problems. Problems might be chosen specifically because they can be solved by heuristic classification.
- *A foundation for characterizing epistemologic adequacy of representation languages* (McCarthy and Hayes, 1969), so that the leverage they provide can be better understood. For example, for classification it is advantageous for a language to provide constructs for representing problem solutions as a network of schemas.
- *A focus for cognitive studies* of human categorization of knowledge and search strategies for retrieval and matching, suggesting principles that might be used in expert programs. Human learning research might similarly focus on the inference structure of heuristic classification.

Finally, it is important to remember that expert systems are programs. Basic computational ideas such as input, output, and sequence, are essential for describing what they do. The methodology of our study has been to ask, "What does the program conclude about? How does it get there from its input?" We characterize the flow of inference, identifying data abstractions, heuristics, implicit models and assumptions, and solution categories along the way. If heuristic programming is to be different from traditional programming, a knowledge-level analysis should always be pursued at least one level deeper than our representations, even if practical constraints prevent making explicit in the implemented program everything that we know. In this way, *knowledge engineering* can be based on sound principles that unite it with studies of cognition and representation.

Acknowledgments

These ideas were originally stimulated by discussions with Denny Brown in our attempt to develop a framework for teaching knowledge engineering. I am grateful to Teknowledge for providing an environment that allowed me to study and analyze a variety of expert systems.

I would also like to thank Tom Dietterich, Steve Hardy, and Peter Szolovits for their suggestions and early encouragement. Discussions with Jim Bennett, John Boose, Keith Butler, Lee Erman, Rick Hayes-Roth, Ramesh Patil, Paul Rosenbloom, Kurt Van Lehn, and Beverly Woolf have also been helpful.

This research has been supported in part by ONR and ARI Contract N00014-79C-0302 and the Josiah Macy, Jr. Foundation. Computational resources have been provided by the SUMEX-

AIM facility (NIH grant RR00785).

References

- Abelson, H., Sussman, G. J. and Sussman, J. *Structure and Interpretation of Computer Programs*. Cambridge: The MIT Press 1985.
- Aiello, N. *A comparative study of control strategies for expert systems: AGE implementation of three variations of PUFF*, in *Proceedings of the National Conference on AI*, pages 1-4, Washington, D.C., August, 1983.
- Aikins J. S. Prototypical knowledge for expert systems. *Artificial Intelligence*, 1983, 20(2), 163-210.
- Amarel, S. *Basic themes and problems in current AI research*, in *Proceedings of the Fourth Annual AIM Workshop*, pages 28-46, June, 1978.
- Bennett, J., Creary, L., Englemore, R., and Melosh, R. *SACON: A knowledge-based consultant for structural analysis*. STAN-CS-78-699 and HPP Memo 78-23, Stanford University, Sept 1978.
- Bobrow, D. G. Qualitative reasoning about physical systems: An introduction. *Artificial Intelligence*, 1984, 24(1-3), 1-5.
- Bobrow, D. G. and Winograd, T. KRL: Another perspective. *Cognitive Science*, 1979, 3, 29-42.
- Boose, J. *Personal construct theory and the transfer of human expertise*, in *Proceedings of the National Conference on AI*, pages 27-33, Austin, TX, August, 1984.
- Borning, A. *Thinglab: A constraint-oriented simulation laboratory*. STAN-CS 79-746, Stanford University, July 1979.
- Brachman, R. J. What's in a concept: Structural foundations for semantic networks. *International Journal of Man-Machine Studies*, 1977, 9, 127-152.
- Brachman, R. J., Fikes, R. E., and Levesque, H. J. KRYPTON: A functional approach to knowledge representation. *IEEE Computers*, 1983, 16(10), 67-73.
- Brown, J. S., Burton, R. R., and de Kleer, J. Pedagogical, natural language, and knowledge engineering techniques in SOPHIE I, II, and III. In D. Sleeman and J. S. Brown (editors), *Intelligent Tutoring Systems*, pages 227-282. Academic Press, London, 1982.
- Bruner, J. S. *The Process of Education*. Cambridge: Harvard University Press 1960.
- Bruner, J. S., Goodnow, J. J., and Austin, G. A. *A Study of Thinking*. New York: John Wiley & Sons, Inc. 1956.
- Buchanan, B. G. and Shortliffe, E. H. *Rule-based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*. Reading: Addison-Wesley Publishing Company 1984.
- Buchanan, B. G., Sutherland, G., and Feigenbaum, E. A. Heuristic dendral: A program for generating explanatory hypotheses in organic chemistry. In B. Meltzer and D. Michie (editors), *Machine Intelligence*, pages 209-254. Edinburgh University Press, Edinburgh, 1969.
- Burton, R. R. Diagnosing bugs in a simple procedural skill. In D. Sleeman and J. S. Brown (editors), *Intelligent Tutoring Systems*, pages 157-183. Academic Press, New York, 1982.

- Chandrasekaran, B. Expert systems: Matching techniques to tasks. In W. Reitman (editor), *AI Applications for Business*, pages 116-132. Ablex Publishing Corp., 1984.
- Chandrasekaran, B. and Mittal, S. Conceptual representation of medical knowledge. In M. Yovits (editor), *Advances in Computers*, pages 217-293. Academic Press, New York, 1983.
- Chi, M. T. H., Feltovich, P. J., Glaser, R. Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 1981, 5, 121-152.
- Clancey, W. J. *The advantages of abstract control knowledge in expert system design*, in *Proceedings of the National Conference on AI*, pages 74-78, Washington, D.C., August, 1983.
- Clancey, W. J. The epistemology of a rule-based expert system: A framework for explanation. *Artificial Intelligence*, 1983, 20(3), 215-251.
- Clancey, W. J. *Acquiring, representing, and evaluating a competence model of diagnosis*. HPP Memo 84-2, Stanford University, February 1984. (To appear in M. Chi, R. Glaser, and M. Farr (Eds.), *Contributions to the Nature of Expertise*, in preparation.).
- Clancey, W. J. *Knowledge acquisition for classification expert systems*, in *Proceedings of ACM Annual Conference*, pages 11-14, October, 1984.
- Clancey, W. J. Details of the revised therapy algorithm. In B. G. Buchanan and E. H. Shortliffe (editors), *Rule-based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*, pages 133-146". Addison-Wesley, 1984.
- Clancey, W. J. Representing control knowledge as abstract tasks and metarules. (To appear in *Computer Expert Systems*, eds. M. J. Coombs and L. Bolc, Springer-Verlag, in preparation).
- Clancey, W. J. and Letsinger, R. NEOMYCIN: Reconfiguring a rule-based expert system for application to teaching. In Clancey, W. J. and Shortliffe, E. H. (editors), *Readings in Medical Artificial Intelligence: The First Decade*, pages 361-381. Addison-Wesley, Reading, 1984.
- Cohen, B. and Murphy, G. L. Models of concepts. *Cognitive Science*, 1984, 8, 27-58.
- Cox, D. J. FLIPP: A method for acquiring and displaying domain knowledge for expert systems. (Center for Creativity, Inc.), Cincinnati, OH.
- Davis R. *Applications of meta-level knowledge to the construction, maintenance, and use of large knowledge bases*. HPP Memo 76-7 and AI Memo 283, Stanford University, July 1976.
- De George, R. T. and De George, F. M. (editors). *The Structuralists: From Marx to Levi-Strauss*. Garden City: Doubleday & Company, Inc., Anchor Books 1972.
- de Kleer, J. Qualitative and quantitative reasoning in classical mechanics. In P. H. Winston and R. H. Brown (editors), *Artificial Intelligence: An MIT Perspective*, pages 9-30. The MIT Press, Cambridge, 1979.
- Engelmore, R. and Terry, A. *Structure and function of the CRYSLIS system*, in *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, pages 250-256, August, 1979.

- Erman, L. D., London, P. E., and Fickas, S. F. *The design and example use of Hearsay-III*, in *Proceedings Seventh International Joint Conference on Artificial Intelligence*, pages 409-415, August, 1981.
- Feigenbaum, E. A. *The art of artificial intelligence: I. Themes and case studies of knowledge engineering*, in *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, pages 1014-1029, August, 1977.
- Flores, C. F. and Winograd, T. *Understanding computers and cognition: A new foundation for design*. Norwood, NJ: Ablex 1985.
- Fox, M. S. and Smith, S. F. ISIS-A knowledge-based system for factory scheduling. *Expert Systems*, 1984, 1(1), 25-49.
- Friedland, P. E. *Knowledge-based experiment design in molecular genetics*. Technical Report STAN-CS-79-771, Stanford University, October 1979.
- Genesereth, M. R. *An overview of meta-level architecture*, in *Proceedings of The National Conference on Artificial Intelligence*, pages 119-124, August, 1983.
- Genesereth, M. R. The use of design descriptions in automated diagnosis. *Artificial Intelligence*, 1984, 24(1-3), 411-436.
- Genesereth, M.R., Greiner, R., Smith, D.E. *MRS Manual*. Heuristic Programming Project Memo HPP-80-24, Stanford University, December 1981.
- Gentner, D. and Stevens, A. (editors). *Mental models*. Hillsdale, NJ: Erlbaum 1983.
- Goldberg, A. and Robson, D. *Smalltalk-80: The Language and its Implementation*. Menlo Park: Addison-Wesley Publishing Company 1983.
- Gomez, F. and Chandrasekaran, B. Knowledge organization and distribution for medical diagnosis. In W. J. Clancey and E. H. Shortliffe (editors), *Readings in Medical Artificial Intelligence: The First Decade*, pages 320-338. Addison-Wesley Publishing Company, Reading, 1984.
- Greeno, J. G. and Simon, H. A. *Problem solving and reasoning*. UPITT/LRDC/ONR/APS 14, University of Pittsburgh, February 1984. To appear in Stevens' Handbook of Experimental Psychology, (Revised Edition). New York: John Wiley & Sons.
- Harmon, P. and King D. *Expert Systems: Artificial Intelligence in Business*. New York: John Wiley & Sons 1985.
- Hart, P. E. *Observations on the development of expert knowledge-based systems*, in *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, pages 1001-1003, August, 1977.
- Hayes, P.J. *In defence of logic*, in *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*, pages 559-565, August, 1977.
- Hayes, P. The logic of frames. In D. Metzger (editor), *Frame Conceptions and Text Understanding*, pages 45-61. de Gruyter, 1979.
- Hayes-Roth, B. and Hayes-Roth, F. A cognitive model of planning. *Cognitive Science*, 1979, 3, 275-310.
- Hayes-Roth, F., Waterman, D., and Lenat, D. (eds.). *Building Expert Systems*. New York:

Addison-Wesley 1983.

- Hewitt, C. Control structure as patterns of passing messages. In P. H. Winston and R. H. Brown (editors), *Artificial Intelligence: An MIT Perspective (Volume 2)*, pages 433-465. The MIT Press, Cambridge, 1979.
- Johnson-Laird, P. N. and Wason, P. C. *Thinking: Readings in Cognitive Science*. Cambridge: Cambridge University Press 1977.
- Kolodner, J. L. *The role of experience in development of expertise*, in *Proceedings of the National Conference on AI*, pages 273-277, Pittsburgh, PA, August, 1982.
- Kolodner, J. L. Maintaining organization in a dynamic long-term memory. *Cognitive Science*, 1983, 7, 243-280.
- Lane, W. G. Input/output processing. In Stone, H. S. (editor), *Introduction to Computer Architecture, 2nd Edition*, chapter 6. Science Research Associates, Inc., Chicago, 1980.
- Lehnert, W., and Wilks, Y. A critical perspective on KRL. *Cognitive Science*, 1979, 3, 1-28.
- Lenat, D. B. and Brown, J. S. Why AM and EURISKO appear to work. *Artificial Intelligence*, 1984, 23(3), 269-294.
- Levesque, H.J. Foundations to a functional approach to knowledge representation. *Artificial Intelligence*, 1984, 23(2), 155-212.
- Mackinlay, J. and Genesereth, M. R. *Expressiveness of languages*, in *Proceedings of the National Conference on Artificial Intelligence*, pages 226-232, August, 1984.
- Martin, W. A. Descriptions and the specialization of concepts. In Winston, P. H. and Brown, R. H. (editor), *Artificial Intelligence: An MIT Perspective*, pages 375-419. The MIT Press, Cambridge, 1979.
- McCarthy, J. and Hayes, P. Some philosophical problems from the standpoint of Artificial Intelligence. In B. Meltzer and D. Michie (editors), *Machine Intelligence 4*, pages 463-502. Edinburgh University Press, 1969.
- McDermott, J. R1: A rule-based configurer of computer systems. *Artificial Intelligence*, 1982, 19(1), 39-88.
- Mervis, C. B. and Rosch, E. Categorization of natural objects. *Annual Review of Psychology*, 1981, 32, 89-115.
- Michalski, R. S. and Stepp, R. E. Learning from observation: conceptual clustering. In R. S. Michalski, R. E. Stepp, and T. M. Mitchell (editors), *Machine Learning*, pages 331-363. Tioga Publishing Company, Palo Alto, 1983.
- Miller, G. A. Practical and lexical knowledge. In E. Rosch and B. B. Lloyd (editors), *Cognition and Categorization*, pages 305-319. Lawrence Erlbaum Associates, Hillsdale, NJ, 1978.
- Mitchell, T. M. *Toward combining empirical and analytical methods for inferring heuristics*. LCSR-TR 27, Laboratory for Computer Science Research, March 1982.
- Neches, R., Swartout, W. R., and Moore, J. *Explainable (and maintainable) expert systems*, in *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 382-389, Los Angeles, CA, August, 1985.

- Newell, A. The knowledge level. *Artificial Intelligence*, 1982, 18(1), 87-127.
- Nii, H. P., Feigenbaum, E. A., Anton, J. J., and Rockmore, A. J. Signal-to-symbol transformation: HASP/SIAP case study. *The AI Magazine*, 1982, 3(2), 23-35.
- Nilsson, N. J. The interplay between theoretical and experimental methods in Artificial Intelligence. *Cognition and Brain Theory*, 1981, 4(1), 69-74.
- Palmer, S. E. Fundamental aspects of cognitive representation. In E. Rosch and B. B. Lloyd (editors), *Cognition and Categorization*, pages 259-303. Lawrence Erlbaum Associates, Hillsdale, NJ, 1978.
- Patil, R. S., Szolovits, P., and Schwartz, W. B. *Causal understanding of patient illness in medical diagnosis*, in *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, pages 893-899, August, 1981.
- Patil, R. S. *Causal representation of patient illness for electrolyte and acid-base diagnosis*. MIT/LCR/TR 267, Massachusetts Institute of Technology, October 1981.
- Pauker, S. G., Gorry, G. A., Kassirer, J. P., and Schwartz, W. B. Toward the simulation of clinical cognition: taking a present illness by computer. *AJM*, 1976, 60, 981-995.
- Pople, H. Heuristic methods for imposing structure on ill-structured problems: the structuring of medical diagnostics. In P. Szolovits (editor), *Artificial Intelligence in Medicine*, pages 119-190. Westview Press, 1982.
- Quillian, M. R. Semantic memory. In M. Minsky (editor), *Semantic Information Processing*. MIT Press, Cambridge, MA, 1968.
- Rich, E. User modeling via stereotypes. *Cognitive Science*, 1979, 3, 355-366.
- Rosch, E. Principles of categorization. In E. Rosch and B. B. Lloyd (editors), *Cognition and Categorization*, pages 27-48. Lawrence Erlbaum Associates, Hillsdale, NJ, 1978.
- Rosch E. and Lloyd B. B. (editors). *Cognition and Categorization*. Hillsdale, NJ: Lawrence Erlbaum Associates 1978.
- Rubin, A. D. *Hypothesis formation and evaluation in medical diagnosis*. Technical Report AI-TR-316, Artificial Intelligence Laboratory, Massachusetts Institute of Technology, January 1975.
- Rumelhart, D. E. and Norman, D. A. *Representation in memory*. Technical Report CHIP-116, Center for Human Information Processing, University of California, June 1983.
- Schank, R. C., and Abelson, R. P. *Scripts, Plans, Goals, and Understanding*. Hillsdale, NJ: Lawrence Erlbaum Associates 1975.
- Schank, R. C. Failure-driven memory. *Cognition and Brain Theory*, 1981, 4(1), 41-60.
- Schmolze, J. G. and Lipkis, T. A. *Classification in the KL-ONE knowledge representation system*, in *Proceedings of the Eighth International Joint Conference on Artificial Intelligence*, pages 330-332, August, 1983.
- Simon, H. A. *The Sciences of the Artificial*. Cambridge: The MIT Press 1969.
- Simon, H. A. The structure of ill structured problems. *Artificial Intelligence*, 1973, 4, 181-201.

- Sowa, J. F. *Conceptual Structures*. Reading, MA: Addison-Wesley 1984.
- Stefik, M. *Planning with constraints*. STAN-CS-80-784 and HPP Memo 80-2, Stanford University, January 1980.
- Stefik, M., Bobrow, D. G., Mittal, S., and Conway, L. Knowledge programming in loops: Report on an experimental course. *The AI Magazine*, 1983, 4(3), 3-13.
- Sussman, G. J. *A Computer Model of Skill Acquisition*. New York: American Elsevier 1975.
- Swartout W. R. *Explaining and justifying in expert consulting programs*, in *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, pages 815-823, August, 1981.
- Szolovits, P. and Pauker, S. G. Categorical and probabilistic reasoning in medical diagnosis. *Artificial Intelligence*, 1978, 11, 115-144.
- van Melle, W. *A domain-independent production rule system for consultation programs*, in *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, pages 923-925, August, 1979.
- VanLehn, K. *Human procedural skill acquisition: Theory, model, and psychological validation*, in *Proceedings of the National Conference on AI*, pages 420-423, Washington, D.C., August, 1983.
- Vemuri, V. *Modeling of Complex Systems: An Introduction*. New York: Academic Press 1978.
- Weiss, S. M. and Kulikowski, C. A. *EXPERT: A system for developing consultation models*, in *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, pages 942-947, August, 1979.
- Weiss, S. M. and Kulikowski, C. A. *A Practical Guide to Designing Expert Systems*. Totowa, NJ: Rowman and Allanheld 1984.
- Weiss, S. M., Kulikowski, C. A., Amarel, S., and Safir, A. A model-based method for computer-aided medical decision making. *Artificial Intelligence*, 1978, 11, 145-172.
- Woods, W. A. What's in a link: Foundations for semantic networks. In D. G. Bobrow and A. Collins (editors), *Representation and Understanding*, pages 35-82. Academic Press, New York, 1975.

Personnel Analysis Division,
AF/MPXA
5C360, The Pentagon
Washington, DC 20330

Air Force Human Resources Lab
AFHRL/MPD
Brooks AFB, TX 78235

AFOSR,
Life Sciences Directorate
Bolling Air Force Base
Washington, DC 20332

Dr. Robert Ahlers
Code N711
Human Factors Laboratory
NAVTRAEQUIPCEN
Orlando, FL 32813

Dr. Ed Aiken
Navy Personnel R&D Center
San Diego, CA 92152

Dr. William E. Alley
AFHRL/MOT
Brooks AFB, TX 78235

Dr. Earl A. Alluisi
HQ, AFHRL (AFSC)
Brooks AFB, TX 78235

Dr. John R. Anderson
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

Dr. Nancy S. Anderson
Department of Psychology
University of Maryland
College Park, MD 20742

Dr. Steve Andriole
Perceptronics, Inc.
21111 Erwin Street
Woodland Hills, CA 91367-3713

Dr. John Annett
University of Warwick
Department of Psychology
Coventry CV4 7AJ
ENGLAND

Dr. Phipps Arable
University of Illinois
Department of Psychology
603 E. Daniel St.
Champaign, IL 61820

Technical Director, ARI
5001 Eisenhower Avenue
Alexandria, VA 22333

Special Assistant for Projects,
CASH(M&RA)
5D800, The Pentagon
Washington, DC 20350

Dr. Michael Atwood
ITT - Programming
1000 Cronoque Lane
Stratford, CT 06497

Dr. Patricia Baggett
University of Colorado
Department of Psychology
Box 345
Boulder, CO 80309

Dr. Eva L. Baker
UCLA Center for the Study
of Evaluation
145 Moore Hall
University of California
Los Angeles, CA 90024

Dr. Meryl S. Baker
Navy Personnel R&D Center
San Diego, CA 92152

Dr. Donald E. Bamber
Code 71
Navy Personnel R & D Center
San Diego, CA 92152

Mr. J. Barber
HQS, Department of the Army
DAPE-ZBR
Washington, DC 20310

Capt. J. Jean Belanger
Training Development Division
Canadian Forces Training System
CFTSHQ, CFB Trenton
Astra, Ontario, K0K
CANADA

CDR Robert J. Biersner, USN
Naval Biodynamics Laboratory
P. O. Box 29407
New Orleans, LA 70189

Dr. Memucha Birenbaum
School of Education
Tel Aviv University
Tel Aviv, Ramat Aviv 69978
ISRAEL

Dr. Werner P. Birke
Personalstamm der Bundeswehr
Kolner Strasse 262
D-5000 Koeln 90
FEDERAL REPUBLIC OF GERMANY

Dr. Gautam Biswas
Department of Computer Science
University of South Carolina
Columbia, SC 29208

Dr. John Black
Yale University
Box 11A, Yale Station
New Haven, CT 06520

Arthur S. Blaiwes
Code N711
Naval Training Equipment Center
Orlando, FL 32813

Dr. Robert Blanchard
Navy Personnel R&D Center
San Diego, CA 92152

Cdt. Arnold Bohrer
Sectie Psychologisch Onderzoek
Rekruterings-En Selectiecentrum
Kwartier Koningen Astrid
Bruijnstraat
1120 Brussels, BELGIUM

Dr. Jeff Bonar
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Gordon H. Bower
Department of Psychology
Stanford University
Stanford, CA 94306

Dr. Richard Braby
NTSC Code 10
Orlando, FL 32751

Dr. Robert Breaux
Code N-095R
NAVTRAEQUIPCEN
Orlando, FL 32813

Dr. John S. Brown
XEROX Palo Alto Research
Center
3333 Coyote Road
Palo Alto, CA 94304

Dr. Bruce Buchanan
Computer Science Department
Stanford University
Stanford, CA 94305

Dr. Patricia A. Butler
NTE Mail Stop 1806
1200 19th St., NW
Washington, DC 20208

Dr. Tom Cafferty
Dept. of Psychology
University of South Carolina
Columbia, SC 29208

Dr. Robert Calfee
School of Education
Stanford University
Stanford, CA 94305

Dr. Jaime Carbonell
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

Mr. James W. Carey
Commandant (G-PTE)
U.S. Coast Guard
2100 Second Street, S.W.
Washington, DC 20593

Dr. Susan Carey
Harvard Graduate School of
Education
337 Gutman Library
Appian Way
Cambridge, MA 02138

Dr. Pat Carpenter
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

Dr. Robert Carroll
NAVOP 01E7
Washington, DC 20370

Dr. Fred Chang
Navy Personnel R&D Center
Code 51
San Diego, CA 92152

Dr. Davida Charney
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Eugene Charniak
Brown University
Computer Science Department
Providence, RI 02912

Dr. Michelene Chi
Learning R & D Center
University of Pittsburgh
3039 O'Hara Street
Pittsburgh, PA 15213

Dr. Susan Chipman
Code 442PT
Office of Naval Research
800 N. Quincy St.
Arlington, VA 22217-5000

Mr. Raymond E. Christal
AFHRL/MCE
Brooks AFB, TX 78235

Dr. Yee-Yeen Chu
Perceptronics, Inc.
21111 Erwin Street
Woodland Hills, CA 91367-3713

Dr. William Clancey
Computer Science Department
Stanford University
Stanford, CA 94306

Chief of Naval Education
and Training
Liaison Office
Air Force Human Resource Laboratory
Operations Training Division
Williams AFB, AZ 85224

Assistant Chief of Staff
for Research, Development,
Test, and Evaluation
Naval Education and
Training Command (N-5)
NAS Pensacola, FL 32508

Dr. Allan M. Collins
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

Dr. Stanley Collyer
Office of Naval Technology
800 N. Quincy Street
Arlington, VA 22217

Dr. Lynn A. Cooper
Learning R&D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

Dr. Meredith P. Crawford
American Psychological Association
Office of Educational Affairs
1200 17th Street, N.W.
Washington, DC 20036

Dr. Hans Crombag
University of Leyden
Education Research Center
Boerhaavelaan 2
2334 EN Leiden
The NETHERLANDS

Dr. Lee Cronbach
16 Laburnum Road
Atherton, CA 94205

Dr. Kenneth B. Cross
Anacapa Sciences, Inc.
P.O. Drawer Q
Santa Barbara, CA 93102

Dr. Mary Cross
Department of Education
Adult Literacy Initiative
Room 4145
400 Maryland Avenue, SW
Washington, DC 20202

CTB/McGraw-Hill Library
2500 Garden Road
Monterey, CA 93940

DR Mike Curran
Office of Naval Research
300 N. Quincy St.
Code 270
Arlington, VA 22217-4000

Brown Gallman
AMRL/LAT
Lowry AFB, CO 80230

Dr. Charles E. Davis
Personnel and Training Research
Office of Naval Research
Code 442PT
800 North Quincy Street
Arlington, VA 22217-4000

Mr. Robert Denton
AFMPC/MPCYPR
Randolph AFB, TX 78150

Mr. Paul DiRenzo
Commandant of the Marine Corps
Code LBC-1
Washington, DC 20380

Dr. R. K. Dismukes
Associate Director for Life Sciences
AFOSR
Bolling AFB
Washington, DC 20332

Defense Technical
Information Center
Cameron Station, Bldg 5
Alexandria, VA 22314
Attn: TC
(12 Copies)

Dr. Thomas M. Duffy
Communications Design Center
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Barbara Eason
Military Educator's
Resource Network
InterAmerica Research Associates
1555 Wilson Blvd
Arlington, VA 22209

Edward E. Eddowes
CNATRA N301
Naval Air Station
Corpus Christi, TX 78419

Dr. John Ellis
Navy Personnel R&D Center
San Diego, CA 92252

Dr. Richard Elster
Deputy Assistant Secretary
of the Navy (Manpower)
OASN (M&RA)
Department of the Navy
Washington, DC 20350-1000

Dr. Susan Embretson
University of Kansas
Psychology Department
Lawrence, KS 66045

Dr. Randy Engle
Department of Psychology
University of South Carolina
Columbia, SC 29208

Lt. Col Rich Entlich
HQ, Department of the Army
OCSA(DACS-DFM)
Washington, DC 20310

Dr. William Epstein
University of Wisconsin
W. J. Brosten Psychology Bldg.
1202 W. Johnson Street
Madison, WI 53706

ERIC Facility-Acquisitions
1873 Puzby Avenue
Bethesda, MD 20014

Dr. K. Anders Ericsson
University of Colorado
Department of Psychology
Boulder, CO 80309

Edward Esty
Department of Education, CERI
MS 40
1200 16th St., NW
Washington, DC 20208

Dr. Beatrice J. Farr
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Marshall J. Farr
2520 North Vernon Street
Arlington, VA 22207

Dr. Pat Federico
C/O 511
NH&DC
San Diego, CA 92152

Dr. Jerome A. Feldman
University of Rochester
Computer Science Department
Rochester, NY 14627

Dr. Paul Feltovich
Southern Illinois University
School of Medicine
Medical Education Department
P.O. Box 3926
Springfield, IL 62708

Mr. Wallace Feurzeig
Educational Technology
Bolt Beranek & Newman
10 Moulton St.
Cambridge, MA 02238

Dr. Craig I. Fields
ARPA
1400 Wilson Blvd.
Arlington, VA 22209

Dr. Gerhard Fischer
University of Colorado
Department of Computer Science
Boulder, CO 80309

Dr. Linda Flower
Carnegie-Mellon University
Department of English
Pittsburgh, PA 15213

Dr. Kenneth D. Forbus
University of Illinois
Department of Computer Science
1304 West Springfield Avenue
Urbana, IL 61801

Dr. Carl H. Frederiksen
McGill University
3700 McTavish Street
Montreal, Quebec H3A 1Y2
CANADA

Dr. John R. Frederiksen
Bolt Beranek & Newman
50 Moulton Street
Cambridge, MA 02138

Dr. Norman Frederiksen
Educational Testing Service
Princeton, NJ 08541

Dr. Alfred R. Fregly
AFCSR/NL
Bolling AFB, DC 20332

Dr. Bob Frey
Commandant (G-P-1/2)
USCG HQ
Washington, DC 20593

Dr. Alinda Friedman
Department of Psychology
University of Alberta
Edmonton, Alberta
CANADA T6G 2E9

Dr. R. Edward Heiselman
Department of Psychology
University of California
Los Angeles, CA 90024

Dr. Michael Genesereth
Stanford University
Computer Science Department
Stanford, CA 94305

Dr. Dedre Gentner
University of Illinois
Department of Psychology
603 E. Daniel St.
Champaign, IL 61820

Dr. Don Gentner
Center for Human
Information Processing
University of California
La Jolla, CA 92093

Dr. Robert Glaser
Learning Research
& Development Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

Dr. Arthur M. Glenberg
University of Wisconsin
W. J. Brogden Psychology Bldg.
1202 W. Johnson Street
Madison, WI 53706

Dr. Marvin D. Glock
13 Stone Hall
Cornell University
Ithaca, NY 14853

Dr. Gene L. Gloye
Office of Naval Research
Detachment
1030 E. Green Street
Pasadena, CA 91106-2485

Dr. Sam Glucksberg
Princeton University
Department of Psychology
Green Hall
Princeton, NJ 08540

Dr. Joseph Goguen
Computer Science Laboratory
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025

Dr. Sherrie Gott
AFHRL/MODJ
Brooks AFB, TX 78235

Jordan Grafman, Ph.D.
Department of Clinical
Investigation
Walter Reed Army Medical Center
6825 Georgia Ave., N.W.
Washington, DC 20307-5001

Dr. Richard H. Granger
Department of Computer Science
University of California, Irvine
Irvine, CA 92717

Dr. Wayne Gray
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. James G. Greeneo
University of California
Berkeley, CA 94720

H. William Greenup
Education Advisor (ECU)
Education Center, NDEEC
Quantico, VA 22134

Dipl. Päd. Michael W. Hubon
Universität Düsseldorf
Erziehungswissenschaftliches
Universitätsstr. 1
D-4000 Düsseldorf 1
WEST GERMANY

Dr. Henry M. Halff
Halff Resources, Inc.
4918 33rd Road, North
Arlington, VA 22207

Dr. Ronald K. Hambleton
Laboratory of Psychometric and
Evaluative Research
University of Massachusetts
Amherst, MA 01003

Dr. Cheryl Hamel
NTEC
Orlando, FL 32813

Dr. Ray Hannapel
Scientific and Engineering
Personnel and Education
National Science Foundation
Washington, DC 20550

Stevan Harnad
Editor, The Behavioral and
Brain Sciences
20 Nassau Street, Suite 240
Princeton, NJ 08540

Mr. William Hartung
PEAM Product Manager
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Wayne Harvey
SRI International
333 Ravenswood Ave.
Room B-S324
Menlo Park, CA 94025

Dr. Reid Hastie
Northwestern University
Department of Psychology
Evanston, IL 60201

Dr. Harold Hawkins
Office of Naval Research
Code 442PT
800 N. Quincy Street
Arlington, VA 22217-5000

Prof. John R. Hayes
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Barbara Hayes-Roth
Department of Computer Science
Stanford University
Stanford, CA 95305

Dr. Frederick Hayes-Roth
Teknowledge
525 University Ave.
Palo Alto, CA 94301

Dr. Joan I. Heller
505 Haddon Road
Oakland, CA 94606

Dr. Geoffrey Hinton
Computer Science Department
Carnegie-Mellon University
Pittsburgh, PA 15213

Dr. Jim Hollan
Code 51
Navy Personnel R & D Center
San Diego, CA 92152

Dr. Melissa Holland
Army Research Institute for the
Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Keith Holyoak
University of Michigan
Human Performance Center
330 Packard Road
Ann Arbor, MI 48109

Prof. Lutz F. Hornke
Universität Düsseldorf
Erziehungswissenschaftliches
Universitätsstr. 1
Düsseldorf 1
WEST GERMANY

Mr. Dick Hoshaw
NAVOP-135
Arlington Annex
Rm 2834
Washington, DC 20350

Dr. Steven Hunka
Department of Education
University of Alberta
Edmonton, Alberta
CANADA

Dr. Earl Hunt
Department of Psychology
University of Washington
Seattle, WA 98105

Dr. Ed Hutchins
Navy Personnel R&D Center
San Diego, CA 92152

Dr. Dillon Inouye
WICAT Education Institute
Provo, UT 84057

Dr. Alice Isen
Department of Psychology
University of Maryland
Catonsville, MD 21228

Dr. Zachary Jacobson
Bureau of Management Consulting
365 Laurier Avenue West
Ottawa, Ontario K1A 0S5
CANADA

Dr. Robert Jannarone
Department of Psychology
University of South Carolina
Columbia, SC 29208

Dr. Claude Janvier
Directeur, CIRADE
Université du Québec à Montréal
Montréal, Québec H3C 3P8
CANADA

COL Dennis W. Jarvi
Commander
AFHRL
Brooks AFB, TX 78235-5601

Margaret Jerome
c/o Dr. Peter Chandler
83, The Drive
Hove
Sussex
UNITED KINGDOM

Dr. Joseph E. Johnson
Assistant Dean for
Graduate Studies
College of Science and Mathematics
University of South Carolina
Columbia, SC 29208

Dr. Douglas H. Jones
Advanced Statistical
Technologies Corporation
10 Trafalgar Court
Lawrenceville, NJ 06143

Col. Dominique Jouslin de Moray
Etat-Major de l'Armee de Terre
Centre de Relations Humaines
3 Avenue Octave Greard
75007 Paris
FRANCE

Dr. Marcel Just
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Milton S. Katz
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Scott Kelso
Haskins Laboratories,
270 Crown Street
New Haven, CT 06510

Dr. Dennis Kibler
University of California
Department of Information
and Computer Science
Irvine, CA 92717

Dr. David Kieras
University of Michigan
Technical Communication
College of Engineering
1223 E. Engineering Building
Ann Arbor, MI 48109

Dr. Peter Kincaid
Training Analysis
& Evaluation Group
Department of the Navy
Orlando, FL 32813

Dr. Walter Kintsch
Department of Psychology
University of Colorado
Campus Box 345
Boulder, CO 80302

Dr. David Klahr
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Mazie Kherr
Program Manager
Training Research Division
HMMRRO
1100 S. Washington
Alexandria, VA 22314

Dr. Janet L. Kolodner
Georgia Institute of Technology
School of Information
& Computer Science
Atlanta, GA 30332

Dr. Stephen Kosslyn
Harvard University
1276 William James Hall
33 Kirkland St.
Cambridge, MA 02138

Dr. Kenneth Kotovskiy
Department of Psychology
Community College of
Allegheny County
800 Allegheny Avenue
Pittsburgh, PA 15213

Dr. David H. Krantz
2 Washington Square Village
Apt. # 150
New York, NY 10012

Dr. Benjamin Kuipers
Department of Mathematics
Dartmouth University
Hanford, NH 03055

Dr. Patrick Kyllonen
AFHRL/MOE
Brooks AFB, TX 78235

Dr. David R. Lambert
Naval Ocean Systems Center
Code 441T
271 Catalina Boulevard
San Diego, CA 92152

Dr. Pat Langley
University of California
Department of Information
and Computer Science
Irvine, CA 92717

M. Diane Langston
Communications Design Center
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Kathleen LaPiana
Naval Health Sciences
Education and Training Command
Naval Medical Command,
National Capital Region
Bethesda, MD 20814-5022

Dr. Jill Larkin
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

Dr. Robert Lawler
Information Sciences, FRL
GTE Laboratories, Inc.
40 Sylvan Road
Waltham, MA 02254

Dr. Paul E. Lehner
PAR Technology Corp.
7926 Jones Branch Drive
Suite 170
McLean, VA 22102

Dr. Alan M. Lesgold
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Alan Leshner
Deputy Division Director
Behavioral and Neural Sciences
National Science Foundation
1800 G Street
Washington, DC 20550

Dr. Jim Levin
University of California
Laboratory for Comparative
Human Cognition
D003A
La Jolla, CA 92093

Dr. Michael Levine
Educational Psychology
210 Education Bldg.
University of Illinois
Champaign, IL 61801

Dr. Charles Lewis
Faculteit Sociale Wetenschappen
Rijksuniversiteit Groningen
Code 30terinestraat 27
97123 Groningen
The NETHERLANDS

Dr. Clayton Lewis
University of Colorado
Department of Computer Science
Campus Box 470
Boulder, CO 80309

Dr. Charlotte Linde
SRI International
333 Ravenswood Avenue
Menlo Park, CA 94025

Dr. Marcia C. Linn
Lawrence Hall of Science
University of California
Berkeley, CA 94720

Dr. Robert Linn
College of Education
University of Illinois
Urbana, IL 61801

Dr. Don Lyon
P. O. Box 44
Higley, AZ 85236

Dr. Jane Malin
Mail Code SR 111
NASA Johnson Space Center
Houston, TX 77058

Dr. William L. Maloy
Chief of Naval Education
and Training
Naval Air Station
Pensacola, FL 32508

Dr. Sandra P. Marshall
Department of Psychology
University of California
Santa Barbara, CA 93106

Dr. Manton M. Matthews
Department of Computer Science
University of South Carolina
Columbia, SC 29208

Dr. Richard E. Mayer
Department of Psychology
University of California
Santa Barbara, CA 93106

Dr. James McBride
Psychological Corporation
c/o Harcourt, Brace,
Javanovich Inc.
1250 West 6th Street
San Diego, CA 92101

Dr. Kathleen McKeown
Columbia University
Department of Computer Science
New York, NY 10027

Dr. Joe McLachlan
Navy Personnel R&D Center
San Diego, CA 92152

Dr. James McMichael
Navy Personnel R&D Center
San Diego, CA 92152

Dr. Barbara Means
Human Resources
Research Organization
1100 South Washington
Alexandria, VA 22314

Dr. Arthur Melmed
U. S. Department of Education
724 Brown
Washington, DC 20208

Dr. Al Meyrowitz
Office of Naval Research
Code 433
500 N. Quincy
Arlington, VA 22217-5000

Dr. Ryszard S. Michalski
University of Illinois
Department of Computer Science
1304 West Springfield Avenue
Urbana, IL 61801

Prof. D. Michie
The Turing Institute
76 North Hanover Street
Glasgow G1 6AD, Scotland
UNITED KINGDOM

Dr. George A. Miller
Department of Psychology
Green Hall
Princeton University
Princeton, NJ 08540

Dr. Lance A. Miller
IBM Thomas J. Watson
Research Center
P.O. Box 218
Yorktown Heights, NY 10598

Dr. Mark Miller
Computer*Thought Corporation
1721 West Plano Parkway
Plano, TX 75075

Dr. Andrew R. Molnar
Scientific and Engineering
Personnel and Education
National Science Foundation
Washington, DC 20550

Dr. William Montague
NPRDC Code 13
San Diego, CA 92152

Dr. Tom Moran
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304

Dr. Melvyn Moy
Navy Personnel R & D Center
San Diego, CA 92152

Dr. Allen Munro
Behavioral Technology
Laboratories - USC
1845 S. Elena Ave., 4th Floor
Redondo Beach, CA 90277

Director,
Decision Support
Systems Division, NMPC
Naval Military Personnel Command
N-164
Washington, DC 20370

Director,
Distribution Department, NMPC
N-4
Washington, DC 20370

Director,
Overseas Duty Support
Program, NMPC
N-62
Washington, DC 20370

Head, HRM Operations Branch,
NMPC
N-62F
Washington, DC 20370

Director,
Recreational Services
Division, NMPC
N-65
Washington, DC 20370

Assistant for Evaluation,
Analysis, and MIS, NMPC
N-6C
Washington, DC 20370

Spec. Asst. for Research, Experi-
mental & Academic Programs,
NPRDC (Code 016)
NAS Memphis (75)
Millington, TN 38054

Director,
Research & Analysis Div.,
NAVFORSTCOM Code 22
4015 Wilson Blvd.
Arlington, VA 22203

Assistant for Long Range
Requirements,
CNO Executive Panel
NAVOP OCK
2000 North Beauregard Street
Alexandria, VA 22311

Assistant for Planning MANTRAPERS
NAVOP 01E6
Washington, DC 20370

Assistant for MPT Research,
Development and Studies
NAVOP 01E7
Washington, DC 20370

Head, Military Compensation
Policy Branch
NAVOP 134
Washington, DC 20370

Head,
Workforce Information Section,
NAVOP 140F
Washington, DC 20370

Head,
Family Support Program Branch,
NAVOP 156
1300 Wilson Blvd., Room 828
Arlington, VA 22209

Head, Economic Analysis Branch,
NAVOP 162
Washington, DC 20370

Head — Manpower, Personnel,
Training, & Reserve Team,
NAVOP 914D
5A578, The Pentagon
Washington, DC 20350

Assistant for Personnel
Logistics Planning,
NAVOP 987H
5D772, The Pentagon
Washington, DC 20350

Leadership Management Education
and Training Project Officer,
Naval Medical Command
Code 05C
Washington, DC 20372

Technical Director,
Navy Health Research Ctr.
P.O. Box 85122
San Diego, CA 92138

Dr. T. Niblett
The Turing Institute
36 North Hanover Street
Glasgow G1 2AD, Scotland
UNITED KINGDOM

Dr. Richard E. Nisbett
University of Michigan
Institute for Social Research
Room 5261
Ann Arbor, MI 48109

Dr. Donald A. Norman
Institute for Cognitive Science
University of California
La Jolla, CA 92093

Director, Training Laboratory,
NPRDC (Code 05)
San Diego, CA 92152

Director, Manpower and Personnel
Laboratory,
NPRDC (Code 06)
San Diego, CA 92152

Director, Human Factors
& Organizational Systems Lab,
NPRDC (Code 07)
San Diego, CA 92152

Fleet Support Office,
NPRDC (Code 301)
San Diego, CA 92152

Library, NPRDC
Code P201L
San Diego, CA 92152

Commanding Officer,
Naval Research Laboratory
Code 2627
Washington, DC 20390

Dr. Harry F. O'Neil, Jr.
Training Research Lab
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Stellan Ohlsson
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

Director, Technology Programs,
Office of Naval Research
Code 200
800 North Quincy Street
Arlington, VA 22217-5000

Director, Research Programs,
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217-5000

Office of Naval Research,
Code 433
800 N. Quincy Street
Arlington, VA 22217-5000

Office of Naval Research,
Code 442
800 N. Quincy St.
Arlington, VA 22217-5000

Office of Naval Research,
Code 442EP
800 N. Quincy Street
Arlington, VA 22217-5000

Office of Naval Research,
Code 442PT
800 N. Quincy Street
Arlington, VA 22217-5000
(6 Copies)

Special Assistant for Marine
Corps Matters,
CMR Code 100M
800 N. Quincy St.
Arlington, VA 22217-5000

Dr. Judith Orasanu
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Jesse Orlansky
Institute for Defense Analyses
1801 N. Beauregard St.
Alexandria, VA 22311

Prof. Seymour Papert
20C-109
Massachusetts Institute
of Technology
Cambridge, MA 02139

Dr. James Paulson
Department of Psychology
Portland State University
P.O. Box 751
Portland, OR 97207

Lt. Col. (Dr.) David Payne
AFHRL
Brooks AFB, TX 78235

Dr. Douglas Pearse
DCIEM
Box 2000
Downsview, Ontario
CANADA

Dr. Robert Penn
NPRDC
San Diego, CA 92152

Dr. Nancy Pennington
University of Chicago
Graduate School of Business
1101 E. 58th St.
Chicago, IL 60637

Military Assistant for Training and
Personnel Technology,
OUSD (R & E)

Room 3D129, The Pentagon
Washington, DC 20301

Dr. Ray Perez
ARI (PERI-II)
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. David M. Perkins
Educational Technology Center
337 Gutman Library
Appian Way
Cambridge, MA 02138

ICDR Frank C. Petho, MSC, USN
CNATRA Code N36, Bldg. 1
NAS
Corpus Christi, TX 78419

Administrative Sciences Department,
Naval Postgraduate School
Monterey, CA 93940

Department of Operations Research,
Naval Postgraduate School
Monterey, CA 93940

Department of Computer Science,
Naval Postgraduate School
Monterey, CA 93940

Dr. Tjeerd Plomp
Twente University of Technology
Department of Education
P.O. Box 217
7500 AE ENSCHEDE
THE NETHERLANDS

Dr. Martha Polson
Department of Psychology
Campus Box 346
University of Colorado
Boulder, CO 80309

Dr. Peter Polson
University of Colorado
Department of Psychology
Boulder, CO 80309

Dr. Steven E. Poltrook
MCC
9430 Research Blvd.
Echelon Bldg #1
Austin, TX 78759-6509

Dr. Harry E. Rople
University of Pittsburgh
Decision Systems Laboratory
1750 Scenic Hall
Pittsburgh, PA 15261

Dr. Joseph Rostka
AFMIL PERI-10
Army Research Institute
5001 Eisenhower Ave.
Alexandria, VA 22333

Dr. Lynne Reder
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. James A. Reggia
University of Maryland
School of Medicine
Department of Neurology
22 South Greene Street
Baltimore, MD 21201

CDR Karen Reider
Naval School of Health Sciences
National Naval Medical Center
Bldg. 141
Washington, DC 20814

Dr. Fred Reif
Physics Department
University of California
Berkeley, CA 94720

Dr. Lauren Resnick
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

Dr. Mary S. Riley
Program in Cognitive Science
Center for Human Information
Processing
University of California
La Jolla, CA 92093

William Rizzo
Code 712 NAVTRAEQUITPCEN
Orlando, FL 32813

Dr. Andrew M. Rose
American Institutes
for Research
1055 Thomas Jefferson St., NW
Washington, DC 20007

Dr. William B. Rouse
Georgia Institute of Technology
School of Industrial & Systems
Engineering
Atlanta, GA 30332

Dr. Donald Rubin
Statistics Department
Science Center, Room 608
1 Oxford Street
Harvard University
Cambridge, MA 02138

Dr. David Rumelhart
Center for Human
Information Processing
Univ. of California
La Jolla, CA 92093

Ms. Riitta Ruotsalainen
General Headquarters
Training Section
Military Psychology Office
PL 910
SF-00101 Helsinki 10, FINLAND

Dr. Michael J. Samet
Perceptronics, Inc
6271 Varrel Avenue
Woodland Hills, CA 91364

Dr. Robert Sasmor
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Roger Schank
Yale University
Computer Science Department
P.O. Box 2159
New Haven, CT 06520

Mrs. Birgitte Schneidelbach
Forsvarets Center for Lederskab
Christiansnavns Voldgade 8
1424 Kobenhavn K
DENMARK

Dr. Walter Schneider
University of Illinois
Psychology Department
603 E. Daniel
Champaign, IL 61820

Dr. Alan H. Schoenfeld
University of California
Department of Education
Berkeley, CA 94720

Dr. Janet Schofield
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Marc Sebrechts
Department of Psychology
Wesleyan University
Middletown, CT 06475

Dr. Judith Segal
Room 819F
NIE
1200 19th Street N.W.
Washington, DC 20208

Dr. Robert J. Seidel
US Army Research Institute
5001 Eisenhower Ave.
Alexandria, VA 22333

Dr. Ramsay W. Selden
NIE
Mail Stop 1241
1200 19th St., NW
Washington, DC 20208

Dr. W. Steve Sellman
OASD(MRA&L)
2B269 The Pentagon
Washington, DC 20301

Dr. Michael G. Shafro
ONR Code 442PT
800 N. Quincy Street
Arlington, VA 22217-5000

Dr. Sylvia A. S. Shafro
National Institute of Education
1200 19th Street
Mail Stop 1806
Washington, DC 20208

Dr. T. E. Sheridan
Dept. of Mechanical Engineering
MIT
Cambridge, MA 02139

Dr. Red Shortliffe
Computer Science Department
Stanford University
Stanford, CA 94305

Dr. Lee Shulman
Stanford University
1340 Carmichael Way
Stanford, CA 94305

Dr. Randall Shumaker
Naval Research Laboratory
Code 7510
4555 Overlook Avenue, S.W.
Washington, DC 20375-5000

Dr. Miriam Shustack
Code 51
Navy Personnel R & D Center
San Diego, CA 92152

Dr. Robert S. Siegler
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Herbert A. Simon
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Zita M. Simutis
Instructional Technology
Systems Area
ARI
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. H. Wallace Sinaiko
Manpower Research
and Advisory Services
Smithsonian Institution
801 North Pitt Street
Alexandria, VA 22314

Dr. Derek Sleeman
Stanford University
School of Education
Stanford, CA 94305

Dr. Charles F. Smith
North Carolina State University
Department of Statistics
Raleigh, NC 27695-8703

Dr. Edward E. Smith
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

Dr. Alfred F. Snode
Senior Scientist
Code 7B
Naval Training Equipment Center
Orlando, FL 32813

Dr. Elliot Soloway
Yale University
Computer Science Department
P.O. Box 2158
New Haven, CT 06520

Dr. Richard Sorensen
Navy Personnel R&D Center
San Diego, CA 92152

Dr. Kathryn T. Spoehr
Brown University
Department of Psychology
Providence, RI 02912

James J. Staszewski
Research Associate
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Marian Stearns
CRI International
333 Ravenswood Ave.
Room B-3724
Menlo Park, CA 94025

Dr. Robert Sternberg
Department of Psychology
Yale University
Box 11A, Yale Station
New Haven, CT 06520

Dr. Albert Stevens
Bolt Beranek & Newman, Inc.
10 Moulton Street
Cambridge, MA 02138

Dr. Paul J. Sticha
Senior Staff Scientist
Training Research Division
HumRRO
1100 S. Washington
Alexandria, VA 22314

Dr. Thomas Sticht
Navy Personnel R&D Center
San Diego, CA 92152

Dr. David Stone
KAJ Software, Inc.
3420 East Shea Blvd.
Suite 161
Phoenix, AZ 85028

Cdr Michael Suman, PD 303
Naval Training Equipment Center
Code N51, Comptroller
Orlando, FL 32813

Dr. Hariharan Swaminathan
Laboratory of Psychometric and
Evaluation Research
School of Education
University of Massachusetts
Amherst, MA 01003

Mr. Brad Simpson
Navy Personnel R&D Center
San Diego, CA 92152

Dr. John Tangney
AFOSR/NL
Bolling AFB, DC 20332

Dr. Kikumi Tatsuoaka
CERL
252 Engineering Research
Laboratory
Urbana, IL 61801

Dr. Martin M. Taylor
DCIEM
Box 2000
Downsview, Ontario
CANADA

Dr. Perry W. Thorndyke
FMC Corporation
Central Engineering Labs
1185 Coleman Avenue, Box 580
Santa Clara, CA 95052

Major Jack Thorpe
DARPA
1400 Wilson Blvd.
Arlington, VA 22209

Dr. Douglas Towne
Behavioral Technology Labs
1845 S. Elena Ave.
Redondo Beach, CA 90277

Dr. Amos Tversky
Stanford University
Dept. of Psychology
Stanford, CA 94305

Dr. James Tweeddale
Technical Director
Navy Personnel R&D Center
San Diego, CA 92152

Dr. Paul Twohig
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Headquarters, U. S. Marine Corps
Code MPI-20
Washington, DC 20380

Dr. Kurt Van Lehn
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304

Dr. Beth Warren
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

Dr. David J. Weiss
N660 Elliott Hall
University of Minnesota
75 E. River Road
Minneapolis, MN 55455

Roger Weissinger-Baylon
Department of Administrative
Sciences
Naval Postgraduate School
Monterey, CA 93940

Dr. Donald Weitzman
MITRE
1820 Dolley Madison Blvd.
MacLean, VA 22102

Dr. Keith T. Wescourt
FMC Corporation
Central Engineering Labs
1185 Coleman Ave., Box 580
Santa Clara, CA 95052

Dr. Douglas Wetzel
Code 12
Navy Personnel R&D Center
San Diego, CA 92152

Dr. Barbara White
Bolt Beranek & Newman, Inc.
10 Moulton Street
Cambridge, MA 02238

Dr. Mike Williams
IntelliGenetics
124 University Avenue
Palo Alto, CA 94301

Dr. Hilda Wing
Army Research Institute
5001 Eisenhower Ave.
Alexandria, VA 22333

Dr. Robert A. Wisher
U.S. Army Institute for the
Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Martin F. Wiskoff
Navy Personnel R & D Center
San Diego, CA 92152

Dr. Merlin C. Wittrock
Graduate School of Education
UCLA
Los Angeles, CA 90024

Mr. John H. Wolfe
Navy Personnel R&D Center
San Diego, CA 92152

Dr. Wallace Wulfeck, III
Navy Personnel R&D Center
San Diego, CA 92152

Dr. Joe Yasutake
AFRL/MLT
Lowry AFB, CO 80030

Dr. Masoud Yazdani
Dept. of Computer Science
University of Exeter
Exeter EX4 4LL
Devon, ENGLAND

Major Frank Yohannan, USMC
Headquarters, Marine Corps
(Code MPI-20)
Washington, DC 20380

Mr. Carl York
System Development Foundation
181 Lytton Avenue
Suite 210
Palo Alto, CA 94301

Dr. Joseph L. Young
Memory & Cognitive
Processes
National Science Foundation
Washington, DC 20550

Dr. Steven Zornetzer
Office of Naval Research
Code 440
800 N. Quincy St.
Arlington, VA 22217-5000

Dr. Michael J. Zyda
Naval Postgraduate School
Code 52CK
Monterey, CA 93943

END

FILMED

1-86

DTIC